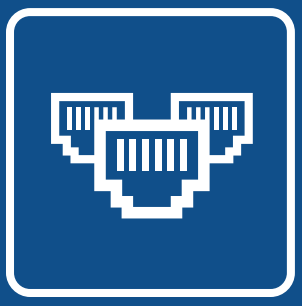

FORTINET
High Performance Network Security

FortiOS™ Handbook - WAN Optimization, Web Cache, Explicit Proxy, and WCCP

VERSION 5.4.1

FORTIOS 5.4 VERSION

FORTINET DOCUMENT LIBRARY

http://docs.fortinet.com

FORTINET VIDEO GUIDE

http://video.fortinet.com

FORTINET BLOG

https://blog.fortinet.com

CUSTOMER SERVICE & SUPPORT

https://support.fortinet.com

http://cookbook.fortinet.com/how-to-work-with-fortinet-support/

FORTIGATE COOKBOOK

http://cookbook.fortinet.com

FORTINET TRAINING SERVICES

http://www.fortinet.com/training

FORTIGUARD CENTER

http://www.fortiguard.com

END USER LICENSE AGREEMENT

http://www.fortinet.com/doc/legal/EULA.pdf

FEEDBACK

Email: techdocs@fortinet.com

Wednesday, June 08, 2016

FortiOS™ Handbook - WAN Optimization, Web Cache, Explicit Proxy, and WCCP

01-541-96996-20160608

TABLE OF CONTENTS

Change Log

Date	Change Description
2016-06-08	Initial release.

Introduction

You can use FortiGate WAN optimization and web caching to improve performance and security of traffic passing between locations on your wide area network (WAN) or from the Internet to your web servers. You can also use the FortiGate unit as an explicit FTP and web proxy server. If your FortiGate unit supports web caching, you can also add web caching to any HTTP sessions including WAN optimization, explicit web proxy and other HTTP sessions.

This document describes how FortiGate WAN optimization, web caching, explicit web proxy, explicit FTP proxy and WCCP work and also describes how to configure these features.

Before you begin

Before you begin to configure WAN optimization, Web caching, explicit proxies or WCCP, take a moment to note the following:

- To use WAN optimization and web caching your FortiGate unit must support these features and not all do. In general your FortiGate unit must include a hard disk to support these features. See FortiGate models that support WAN optimization on page 11. Most FortiGate units support the explicit web and FTP proxies.
- To be able to configure WAN optimization and web caching from the web manager you should begin by going to **System > Feature Select** and turning on **WAN Opt. & Cache**.
- To be able to configure the Explicit Web and FTP proxies from the web manager you should begin by going to **System > Feature Select** and turning on **Explicit Proxy**.
- If you enable virtual domains (VDOMs) on the FortiGate unit, WAN optimization, web caching, and the explicit web and FTP proxies are available separately for each VDOM.
- This guide is based on the assumption that you are a FortiGate administrator. It is not intended for others who may also use the FortiGate unit, such as FortiClient administrators or end users.
- FortiGate WAN optimization is proprietary to Fortinet. FortiGate WAN optimization is compatible only with FortiClient WAN optimization, and will not work with other vendors' WAN optimization or acceleration features.
- FortiGate web caching, explicit web and FTP proxies, and WCCP support known standards for these features. See the appropriate chapters of this document for details.

At this stage, the following installation and configuration conditions are assumed:

- For WAN optimization you have already successfully installed two or more FortiGate units at various locations across your WAN.
- For web caching, the explicit proxies and WCCP you have already successfully installed one or more FortiGate units on your network.
- You have administrative access to the web-based manager and/or CLI.
- The FortiGate units are integrated into your WAN or other networks
- The operation mode has been configured.
- The system time, DNS settings, administrator password, and network interfaces have been configured.
- Firmware, FortiGuard Antivirus and FortiGuard Antispam updates are completed.
- You Fortinet products have been registered. Register your Fortinet products at the Fortinet Technical Support web site, https://support.fortinet.com.

FortiGate models that support WAN optimization

WAN optimization is available on FortiGate models with internal storage that also support SSL acceleration. Internal storage includes high-capacity internal hard disks, AMC hard disk modules, FortiGate Storage Modules (FSMs) or over 4 Gbytes of internal flash storage. All of these storage locations can provide similar web caching and byte caching performance. If you add more than one storage location (for example, by creating multiple partitions on a storage device, by using more than one FSM, or by using an FSM and AMC hard disk in the same FortiGate unit) you can configure different storage locations for web caching and byte caching.

Toggling Disk Usage for logging or wan-opt

Both logging and WAN Optimization use hard disk space to save data. For FortiOS 5.4 you cannot use the same hard disk for WAN Optimization and logging.

- If the FortiGate has one hard disk, then it can be used for either disk logging or WAN optimization, but not both. By default, the hard disk is used for disk logging.
- If the FortiGate has two hard disks, then one disk is always used for disk logging and the other disk is always used for WAN optimization.

On the FortiGate, go to **System > Advanced > Disk Settings** to switch between **Local Log** and **WAN Optimization**.

You can also change disk usage from the CLI using the following command:

```
configure system global
   set disk-usage {log | wanopt}
end
```

The Toggle Disk Usage feature is supported on all new "E" Series models, while support for "D" Series models may vary.

Please refer to the Feature Platform Matrix for more information.

Changing the disk setting formats the disk, erases current data stored on the disk and disables either disk logging or WAN Optimization.

You can configure WAN Optimization from the CLI or the GUI. To configure WAN Optimization from the GUI you must go to **System > Feature Select** and turn on WAN Optimization.

Remote logging (including logging to FortiAnalyzer and remote Syslog servers) is not affected by using the single local hard disk for WAN Optimization.

Enabling WAN Optimization affects more than just disk logging

In addition to affecting WAN Optimization, the following table shows other features affected by the FortiGate disk configuration.

Features affected by Disk Usage as per the number of internal hard disks on the FortiGate

Feature	Logging Only (1 hard disk)	WAN Opt. Only (1 hard disk)	Logging & WAN Opt. (2 hard disks)
Logging	Supported	Not supported	Supported
Report/Historical FortiView	Supported	Not supported	Supported
Firewall Packet Capture (Policy Capture and Interface Capture)	Supported	Not supported	Supported
AV Quarantine	Supported	Not supported	Supported
IPS Packet Capture	Supported.	Not supported	Supported
DLP Archive	Supported	Not supported	Supported
Sandbox DB & Results	FortiSandbox database and results are also stored on disk, but will not be affected by this feature.		

Distributing WAN optimization, explicit proxy, and web caching to multiple CPU Cores

By default WAN optimization, explicit proxy and web caching is handled by half of the CPU cores in a FortiGate unit. For example, if your FortiGate unit has 4 CPU cores, by default two will be used for WAN optimization, explicit proxy and web caching. You can use the following command to change the number of CPU cores that are used.

```
config system global
   set wad-worker-count <number>
end
```

The value for `<number>` can be between 1 and the total number of CPU cores in your FortiGate unit. Adding more cores may enhance WAN optimization, explicit proxy and web caching performance and reduce the performance of other FortiGate systems.

How this chapter is organized

This FortiOS Handbook chapter describes how to implement WAN optimization, web caching and the web proxy on supported FortiGate units.

The FortiOS Handbook chapter contains the following sections:

- Example network topologies provides an overview of FortiGate WAN optimization best practices and technologies and some of the concepts and rules for using them. We recommend that you begin with this chapter before attempting to configure your FortiGate unit to use WAN optimization.
- Configuring WAN optimization provides basic configuration for WAN optimization rules, including adding rules, organizing rules in the rule list and using WAN optimization addresses. This chapter also explains how WAN optimization accepts sessions, as well as how and when you can apply security profiles to WAN optimization traffic.
- Peers and authentication groups describes how to use WAN optimization peers and authentication groups to control access to WAN optimization tunnels.
- Configuration examples describes basic active-passive and peer-to-peer WAN optimization configuration examples. This chapter is a good place to start learning how to put an actual WAN optimization network together.
- Web caching and SSL offloading describes how web caching works to cache HTTP and HTTPS, how to use SSL offloading to improved performance of HTTPS websites, and includes web caching configuration examples.
- FortiClient WAN optimization describes how FortiGate and FortiClient WAN optimization work together and includes an example configuration.
- The FortiGate explicit web proxy describes how to configure the FortiGate explicit web proxy, how users connect to the explicit web proxy, and how to add web caching to the explicit web proxy.
- The FortiGate explicit FTP proxy describes how to configure the FortiGate explicit FTP proxy and how users connect to the explicit FTP proxy.
- FortiGate WCCP describes FortiGate WCCP and how to configure WCCP and the WCCP client.
- Diagnose commands describes get and diagnose commands available for troubleshooting WAN optimization, web cache, and WCCP.

What's new in FortiOS 5.4

Toggle Disk Usage for logging or wan-opt (290892)

Both logging and WAN Optimization use hard disk space to save data. For FortiOS 5.4 you cannot use the same hard disk for WAN Optimization and logging.

- If the FortiGate has one hard disk, then it can be used for either disk logging or WAN optimization, but not both. By default, the hard disk is used for disk logging.
- If the FortiGate has two hard disks, then one disk is always used for disk logging and the other disk is always used for WAN optimization.

On the FortiGate, go to **System > Advanced > Disk Settings** to switch between **Local Log** and **WAN Optimization**.

You can also change disk usage from the CLI using the following command:

```
configure system global
   set disk-usage {log | wanopt}
end
```

The Toggle Disk Usage feature is supported on all new "E" Series models, while support for "D" Series models may vary.

Please refer to the Feature Platform Matrix for more information.

Changing the disk setting formats the disk, erases current data stored on the disk and disables either disk logging or WAN Optimization.

You can configure WAN Optimization from the CLI or the GUI. To configure WAN Optimization from the GUI you must go to **System > Feature Select** and turn on WAN Optimization.

Remote logging (including logging to FortiAnalyzer and remote Syslog servers) is not affected by using the single local hard disk for WAN Optimization.

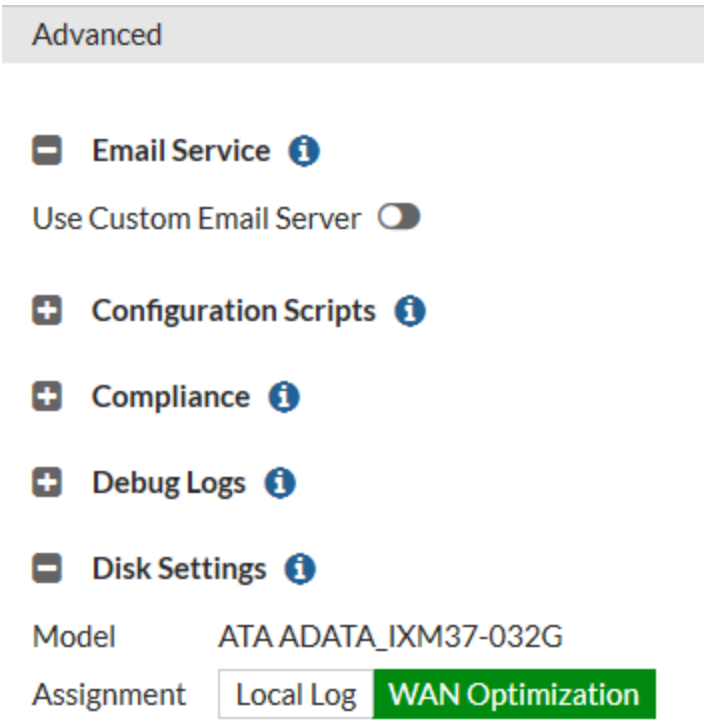

Enabling WAN Optimization affects more than just disk logging

In addition to affecting WAN Optimization, the following table shows other features affected by the FortiGate disk configuration.

Features affected by Disk Usage as per the number of internal hard disks on the FortiGate

Feature	Logging Only (1 hard disk)	WAN Opt. Only (1 hard disk)	Logging & WAN Opt. (2 hard disks)
Logging	Supported	Not supported	Supported
Report/Historical FortiView	Supported	Not supported	Supported
Firewall Packet Capture (Policy Capture and Interface Capture)	Supported	Not supported	Supported
AV Quarantine	Supported	Not supported	Supported
IPS Packet Capture	Supported.	Not supported	Supported
DLP Archive	Supported	Not supported	Supported

Feature	Logging Only (1 hard disk)	WAN Opt. Only (1 hard disk)	Logging & WAN Opt. (2 hard disks)
Sandbox DB & Results	FortiSandbox database and results are also stored on disk, but will not be affected by this feature.		

MAPI AV scanning is supported over WAN Optimization (267975)

AV works on MAPI when WAN Optimization is used.

New explicit proxy features

The following section describes new explicit web proxy features added to FortiOS 5.4.0 and FortiOS 5.4.1.

FortiOS 5.4.1

These features first appeared in FortiOS 5.4.1.

Support Kerberos and NTLM authentication (370489)

FortiGate now recognizes the client's authentication method from the token and selects the correct authentication scheme to authenticate successfully.

CLI syntax

```
config firewall explicit-proxy-policy
   edit <example>
      set active-auth-method [ntlm | basic | digest | negotiate | none]
   end
```

Explicit Web Proxy WISP support improvements (309388 309236)

The following Explicit Web Proxy WISP CLI syntax has been changed and added:

- Changed web-proxy wisp to table object and added outgoing-ip.

CLI syntax

```
config web-proxy
   set server-ip // WISP server IP address
   set server-port // WISP server port (1 - 65535)
```

- In the web filter profile, added WISP servers and WISP algorithm.

CLI syntax

```
config webfilter profile
edit <example>
```

```
set wisp-servers // WISP servers
set wisp-algorithm // WISP server selection algorithm
```

Improvements to explicit web proxy policy page (305817)

Explicit proxy URL categories show description next to their numerical values in the CLI. Also, all categories for **URL Category** are available in the GUI.

Explicit web proxy Kerberos authentication support (297503)

The following web proxy Kerberos authentication CLI syntax has been added:

CLI syntax

```
config user krb-keytab
edit <example>
set principal // Kerberos service principal
set ldap-server // LDAP server name
set keytab // base64 coded keytab
```

Explicit proxy, Web Caching, and WAN Optimization are not supported for Flow-based VDOMs (274748)

Explicit proxy, web caching, and WAN optimization have been removed from the GUI in a Flow-based VDOM.

Explicit proxy support for base64 encoded X-Authenticated-Groups and X-Authenticated-User HTTP headers (356979)

Data for http header-names X-Authenticated-Groups and X-Authenticated-User are decoded before further processing.

FortiOS 5.4.0

These features first appeared in FortiOS 5.4.0.

New explicit proxy firewall address types (284753)

New explicit proxy firewall address types improve granularity over header matching for explicit web proxy policies. You can enable this option using the **Show in Address List** button on the Address and Address Group New/Edit forms under **Policy & Objects > Addresses**.

The following new address types have been added:

- **URL Pattern** - destination address
- **Host Regex Match** - destination address
- **URL Category** - destination address (URL filtering)
- **HTTP Method** - source address
- **User Agent** - source address

- **HTTP Header** - source address
- **Advanced (Source)** - source address (combines User Agent, HTTP Method, and HTTP Header)
- **Advanced (Destination)** - destination address (combines Host Regex Match and URL Category)

Disclaimer messages can be added to explicit proxy policies (273208)

Disclaimer options are now available for each explicit proxy policy or split policy of ID-based policy. This feature allows you to create user exceptions for specific URL categories (including warning messages) based on user groups.

The **Disclaimer Options** are configured under **Policy & Objects > Explicit Proxy Policy**. You can also configure a disclaimer for each Authentication Rule by setting **Action** to **Authenticate**.

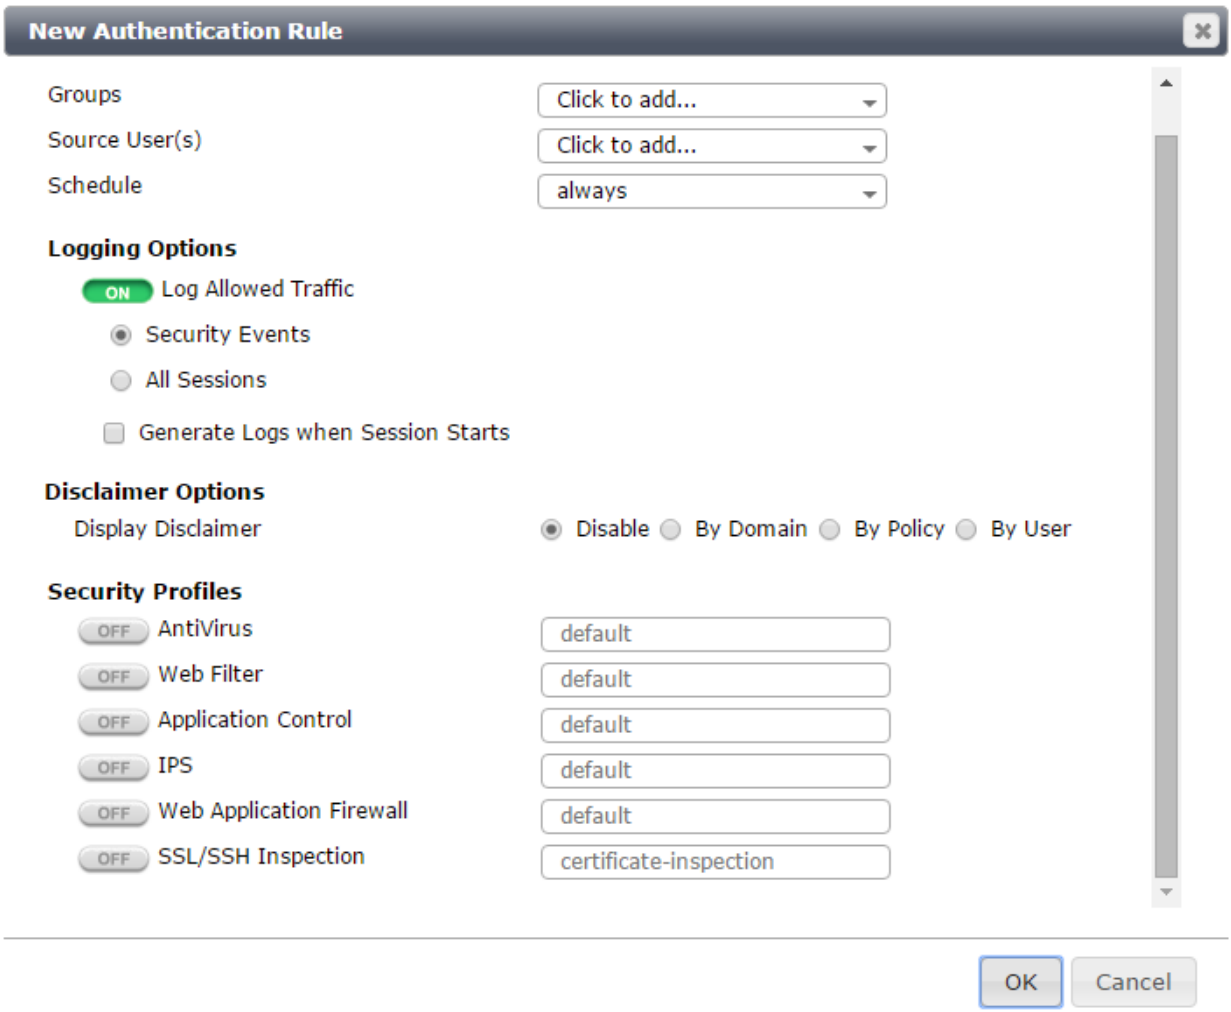

Disclaimer explanations

- **Disable:** No disclaimer (default setting).
- **By Domain:** The disclaimer will be displayed on different domains. The explicit web proxy will check the referring header to mitigate the javascript/css/images/video/etc page.

- **By Policy:** The disclaimer will be displayed ifa the HTTP request matches a different explicit firewall policy.
- **By User:** The disclaimer will be displayed when a new user logs on.

Firewall virtual IPs (VIPs) can be used with Explicit Proxy policies (234974)

The explicit web-proxy will now accept VIP addresses for destination address. If an external IP matches a VIP policy, the IP is changed to the mapped-IP of the VIP.

Implement Botnet features for explicit policy (259580)

The option `scan-botnet-connections` has been added to the firewall explicit proxy policy.

Syntax:

```
config firewall explicit-proxy-policy
   edit <policyid>
      set scan-botnet-connections [disable/block/monitor]
end
```

where:

`disable` means do not scan connections to botnet servers.

`block` means block connections to botnet servers.

`monitor` means log connections to botnet servers.

Add HTTP.REFERRER URL to web filter logs (260538)

Added support for the referrer field in the HTTP header on webfilter log, this field along with others in the HTTP header are very useful in heuristic analysis /search for malware infested hosts.

Adding guest management to explicit web proxy (247566)

Allow user group with type ***Guest*** to be referenced in explicit-proxy-policy.

Example network topologies

FortiGate WAN optimization consists of a number of techniques that you can apply to improve the efficiency of communication across your WAN. These techniques include protocol optimization, byte caching, web caching, SSL offloading, and secure tunneling. Protocol optimization can improve the efficiency of traffic that uses the CIFS, FTP, HTTP, or MAPI protocol, as well as general TCP traffic. Byte caching caches files and other data on FortiGate units to reduce the amount of data transmitted across the WAN. Web caching stores web pages on FortiGate units to reduce latency and delays between the WAN and web servers. SSL offloading offloads SSL decryption and encryption from web servers onto FortiGate SSL acceleration hardware. Secure tunneling secures traffic as it crosses the WAN.

You can apply different combinations of these WAN optimization techniques to a single traffic stream depending on the traffic type. For example, you can apply byte caching and secure tunneling to any TCP traffic. For HTTP and HTTPS traffic, you can also apply protocol optimization and web caching.

You can configure a FortiGate unit to be an explicit web proxy server for both IPv4 and IPv6 traffic and an explicit FTP proxy server. Users on your internal network can browse the Internet through the explicit web proxy server or connect to FTP servers through the explicit FTP proxy server. You can also configure these proxies to protect access to web or FTP servers behind the FortiGate unit using a reverse proxy configuration.

Web caching can be applied to any HTTP or HTTPS traffic, this includes normal traffic accepted by a security policy, explicit web proxy traffic, and WAN optimization traffic.

You can also configure a FortiGate unit to operate as a Web Cache Communication Protocol (WCCP) client or server. WCCP provides the ability to offload web caching to one or more redundant web caching servers.

FortiGate units can also apply security profiles to traffic as part of a WAN optimization, explicit web proxy, explicit FTP proxy, web cache and WCCP configuration. Security policies that include any of these options can also include settings to apply all forms of security profiles supported by your FortiGate unit.

Basic WAN optimization topology

The basic FortiGate WAN optimization topology consists of two FortiGate units operating as WAN optimization peers intercepting and optimizing traffic crossing the WAN between the private networks.

Security device and WAN optimization topology

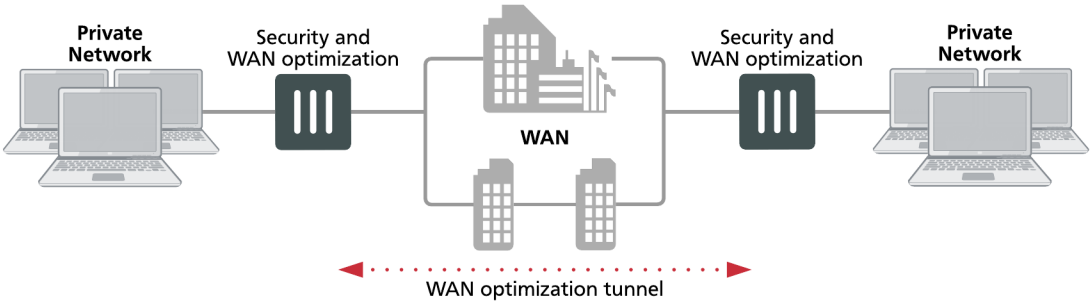

FortiGate units can be deployed as security devices that protect private networks connected to the WAN and also perform WAN optimization. In this configuration, the FortiGate units are configured as typical security devices for the private networks and are also configured for WAN optimization. The WAN optimization configuration intercepts traffic to be optimized as it passes through the FortiGate unit and uses a WAN optimization tunnel with another FortiGate unit to optimize the traffic that crosses the WAN.

You can also deploy WAN optimization on single-purpose FortiGate units that only perform WAN optimization. In the out of path WAN optimization topology shown below, FortiGate units are located on the WAN outside of the private networks. You can also install the WAN optimization FortiGate units behind the security devices on the private networks.

The WAN optimization configuration is the same for FortiGate units deployed as security devices and for single-purpose WAN optimization FortiGate units. The only differences would result from the different network topologies.

Out-of-path WAN Optimization topology

In an out-of-path topology, one or both of the FortiGate units configured for WAN optimization are not directly in the main data path. Instead, the out-of-path FortiGate unit is connected to a device on the data path, and the device is configured to redirect sessions to be optimized to the out-of-path FortiGate unit.

Single-purpose WAN optimization topology

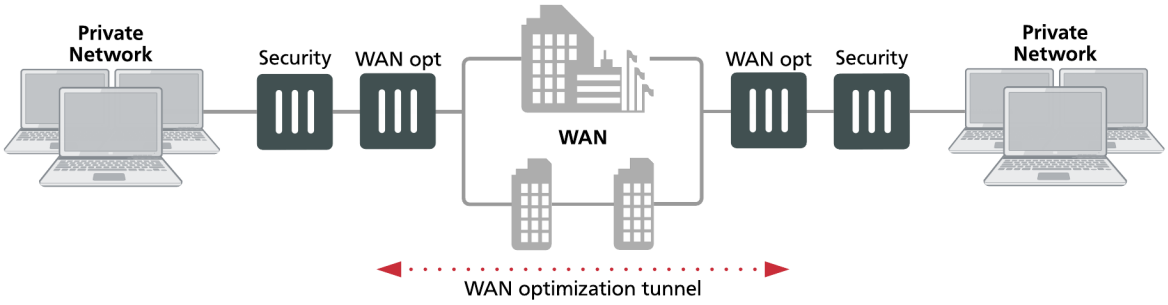

The following out-of-path FortiGate units are configured for WAN optimization and connected directly to FortiGate units in the data path. The FortiGate units in the data path use a method such as policy routing to redirect traffic to be optimized to the out-of-path FortiGate units. The out-of-path FortiGate units establish a WAN optimization tunnel between each other and optimize the redirected traffic.

Out-of-path WAN optimization

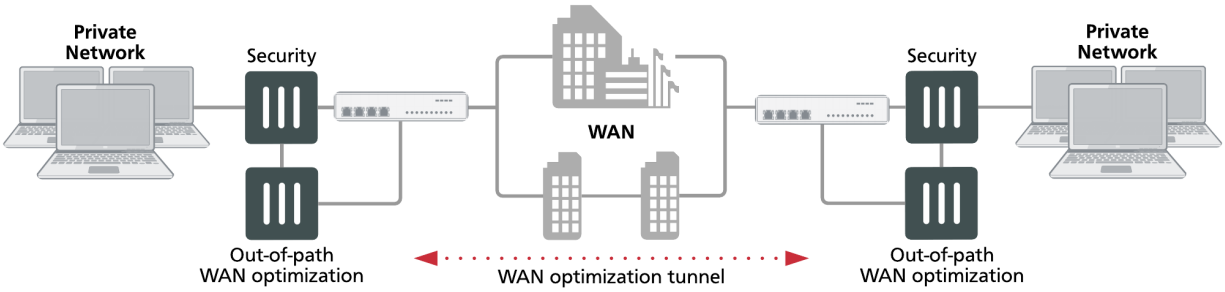

One of the benefits of out-of-path WAN optimization is that out-of-path FortiGate units only perform WAN optimization and do not have to process other traffic. An in-path FortiGate unit configured for WAN optimization also has to process other non-optimized traffic on the data path.

The out-of-path FortiGate units can operate in NAT/Route or Transparent mode.

Other out-of-path topologies are also possible. For example, you can install the out-of-path FortiGate units on the private networks instead of on the WAN. Also, the out-of-path FortiGate units can have one connection to the network instead of two. In a one-arm configuration such as this, security policies and routing have to be configured to send the WAN optimization tunnel out the same interface as the one that received the traffic.

Topology for multiple networks

As shown in below, you can create multiple WAN optimization configurations between many private networks. Whenever WAN optimization occurs, it is always between two FortiGate units, but you can configure any FortiGate unit to perform WAN optimization with any of the other FortiGate units that are part of your WAN.

WAN optimization among multiple networks

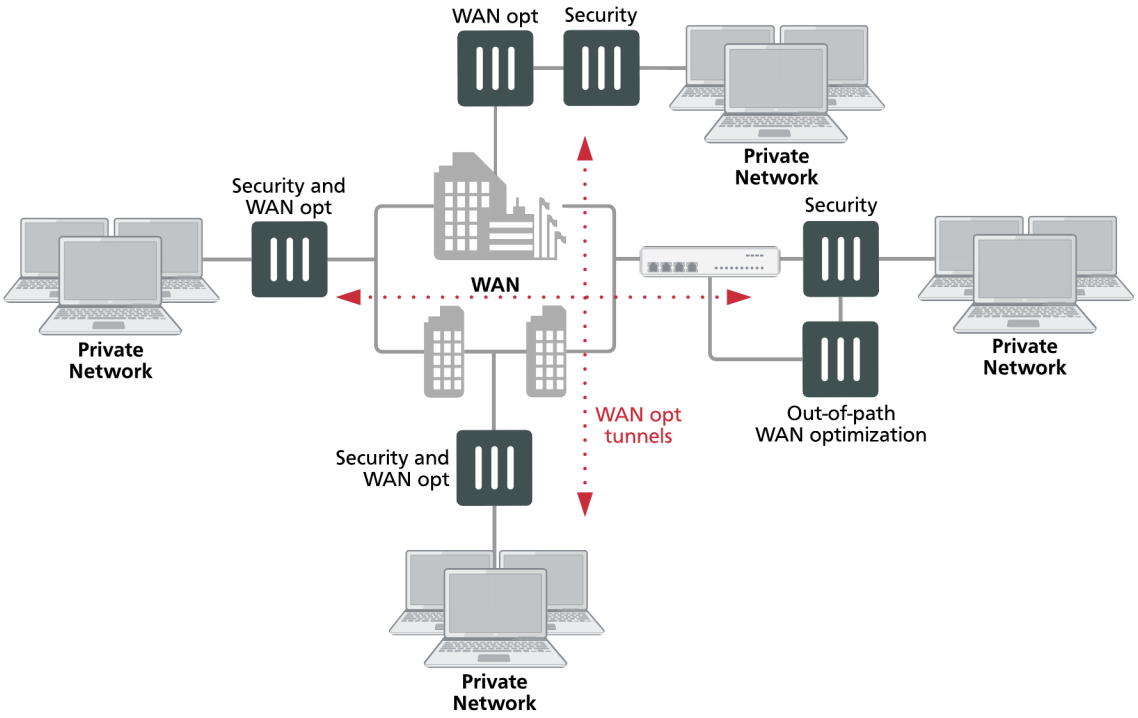

You can also configure WAN optimization between FortiGate units with different roles on the WAN. FortiGate units configured as security devices and for WAN optimization can perform WAN optimization as if they are single-purpose FortiGate units just configured for WAN optimization.

WAN optimization with web caching

You can add web caching to a WAN optimization topology when users on a private network communicate with web servers located across the WAN on another private network.

WAN optimization with web caching topology

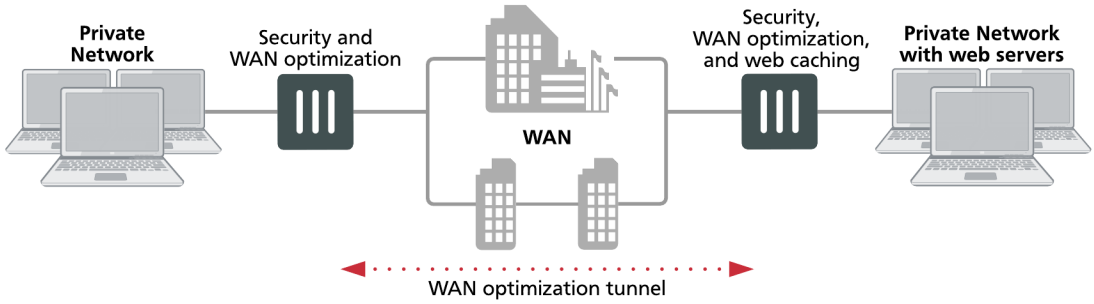

The topology above is the same as that shown in Basic WAN optimization topology on page 20 with the addition of web caching to the FortiGate unit in front of the private network that includes the web servers. You can also add web caching to the FortiGate unit that is protecting the private network. In a similar way, you can add web caching to any WAN Optimization topology.

WAN optimization and web caching with FortiClient peers

FortiClient WAN optimization works with FortiGate WAN optimization to accelerate remote user access to the private networks behind FortiGate units. The FortiClient application requires a simple WAN optimization configuration to automatically detect if WAN optimization is enabled on the FortiGate unit. Once WAN optimization is enabled, the FortiClient application transparently makes use of the WAN optimization and web caching features available.

FortiClient WAN optimization topology

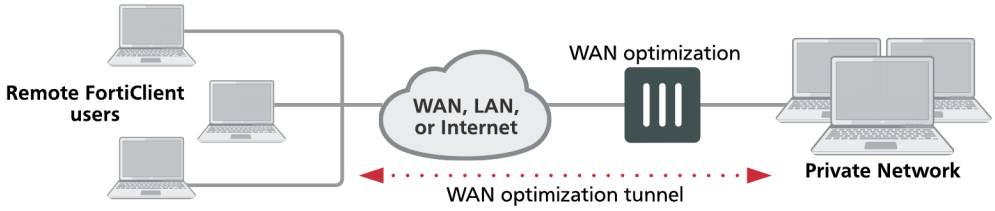

Explicit Web proxy topologies

You can configure a FortiGate unit to be an explicit web proxy server for Internet web browsing of IPv4 and IPv6 web traffic. To use the explicit web proxy, users must add the IP address of the FortiGate interface configured for the explicit web proxy to their web browser proxy configuration.

Explicit web proxy topology

If the FortiGate unit supports web caching, you can also add web caching to the security policy that accepts explicit web proxy sessions The FortiGate unit then caches Internet web pages on a hard disk to improve web browsing performance.

Explicit web proxy with web caching topology

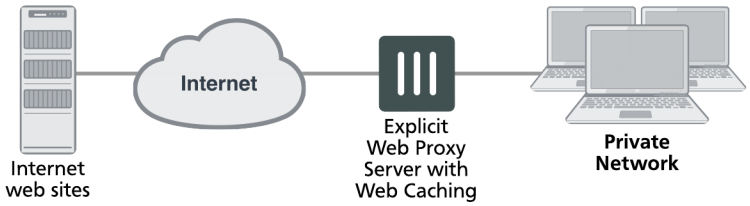

Explicit FTP proxy topologies

You can configure a FortiGate unit to be an explicit FTP proxy server for FTP users. To use the explicit web proxy, FTP users must connect to and authenticate with the explicit FTP proxy before connecting to an FTP server.

Explicit FTP proxy topology

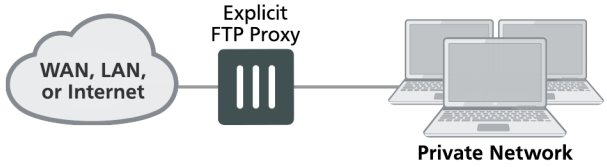

You can also configure reverse explicit FTP proxy. In this configuration, users on the Internet connect to the explicit web proxy before connecting to an FTP server installed behind a FortiGate unit.

Reverse explicit FTP proxy topology

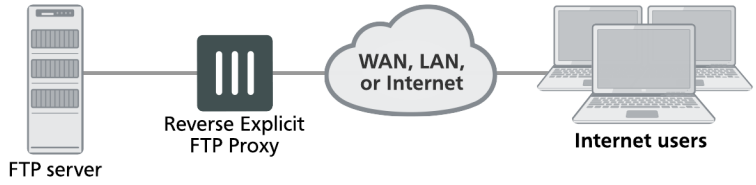

Web caching topologies

FortiGate web caching can be added to any security policy and any HTTP or HTTPS traffic accepted by that security policy can be cached on the FortiGate unit hard disk. This includes WAN optimization and explicit web proxy traffic. The network topologies for these scenarios are very similar. They involved a FortiGate unit installed between users and web servers with web caching enabled.

A typical web-caching topology includes one FortiGate unit that acts as a web cache server. Web caching is enabled in a security policy and the FortiGate unit intercepts web page requests accepted by the security policy, requests web pages from the web servers, caches the web page contents, and returns the web page contents to the users. When the FortiGate unit intercepts subsequent requests for cached web pages, the FortiGate unit contacts the destination web server just to check for changes.

Web caching topology

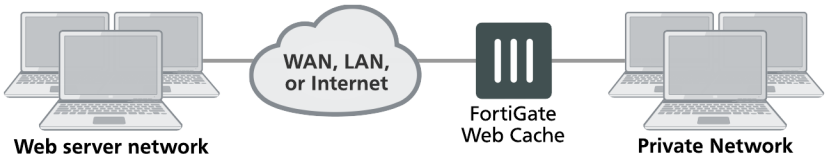

You can also configure reverse proxy web-caching. In this configuration, users on the Internet browse to a web server installed behind a FortiGate unit. The FortiGate unit intercepts the web traffic (HTTP and HTTPS) and caches pages from the web server. Reverse proxy web caching on the FortiGate unit reduces the number of requests that the web server must handle, leaving it free to process new requests that it has not serviced before.

Reverse proxy web caching topology

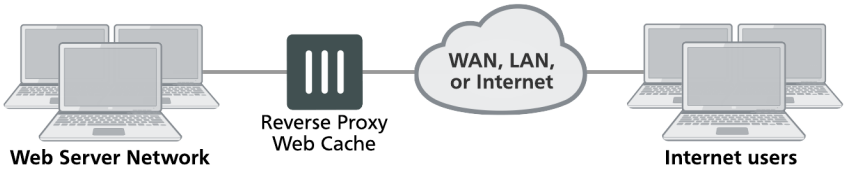

WCCP topologies

You can operate a FortiGate unit as a Web Cache Communication Protocol (WCCP) router or cache engine. As a router, the FortiGate unit intercepts web browsing requests from client web browsers and forwards them to a WCCP cache engine. The cache engine returns the required cached content to the client web browser. If the cache server does not have the required content it accesses the content, caches it and returns the content to the client web browser.

WCCP topology

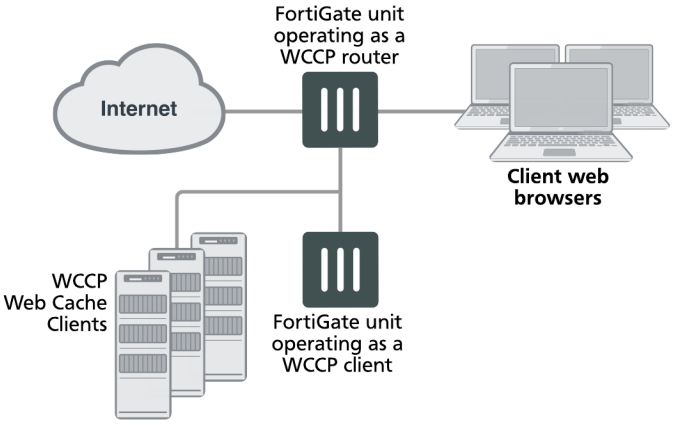

FortiGate units can also operate as WCCP cache servers, communicating with WCCP routers, caching web content and providing it to client web browsers as required.

WCCP is transparent to client web browsers. The web browsers do not have to be configured to use a web proxy.

Configuring WAN optimization

This chapter describes FortiGate WAN optimization client server architecture and other concepts you need to understand to be able to configure FortiGate WAN optimization.

Client/server architecture

Traffic across a WAN typically consists of clients on a client network communicating across a WAN with a remote server network. The clients do this by starting communication sessions from the client network to the server network. These communication sessions can be open text over the WAN or they can be encrypted by SSL VPN or IPsec VPN.

To optimize these sessions, you can add **WAN optimization security policies** to the **client-side FortiGate unit** to accept sessions from the client network that are destined for the server network. The client-side FortiGate unit is located between the client network and the WAN. WAN optimization security policies include **WAN optimization profiles** that control how the traffic is optimized.

The client-side FortiGate unit must also include the IP address of the **server-side FortiGate unit** in its WAN optimization **peer** configuration. The server-side FortiGate unit is located between the server network and the WAN, The peer configuration allows the client-side FortiGate unit to find the server-side FortiGate unit and attempt to establish a WAN optimization **tunnel** with it.

For the server-side FortiGate unit you must add a security policy with **wanopt** as the **Incoming Interface**. This security policy allows the FortiGate unit to accept WAN optimization sessions from the client-side FortiGate unit. For the server-side FortiGate unit to accept a WAN optimization connection it must have the client-side FortiGate unit in its WAN optimization peer configuration.

WAN optimization profiles are only added to the client-side WAN optimization security policy. The server-side FortiGate unit employs the WAN optimization settings set in the WAN optimization profile on the client-side FortiGate unit.

Client/server architecture

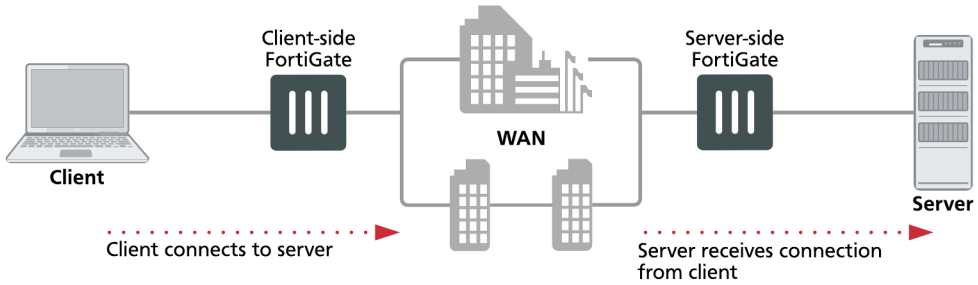

When both peers are identified the FortiGate units attempt to establish a WAN optimization **tunnel** between them. WAN optimization tunnels use port 7810. All optimized data flowing across the WAN between the client-side and server-side FortiGate units use this tunnel. WAN optimization tunnels can be encrypted use SSL encryption to keep the data in the tunnel secure.

Any traffic can be sent through a WAN optimization tunnel. This includes SSL and IPsec VPN traffic. However, instead of configuring SSL or IPsec VPN for this communication you can add SSL encryption using the WAN optimization tunnel.

In addition to basic identification by peer host ID and IP address you can configure WAN optimization **authentication** using certificates and pre-shared keys to improve security. You can also configure FortiGate units involved in WAN optimization to accept connections from any identified peer or restrict connections to specific peers.

The FortiClient application can act in the same manner as a client-side FortiGate unit to optimize traffic between a computer running FortiClient and a FortiGate unit.

WAN optimization peers

The client-side and server-side FortiGate units are called WAN optimization peers because all of the FortiGate units in a WAN optimization network have the same peer relationship with each other. The client and server roles just relate to how a session is started. Any FortiGate unit configured for WAN optimization can be a client-side and a server-side FortiGate unit at the same time, depending on the direction of the traffic. Client-side FortiGate units initiate WAN optimization sessions and server-side FortiGate units respond to the session requests. Any FortiGate unit can simultaneously be a client-side FortiGate unit for some sessions and a server-side FortiGate unit for others.

WAN optimization peer and tunnel architecture

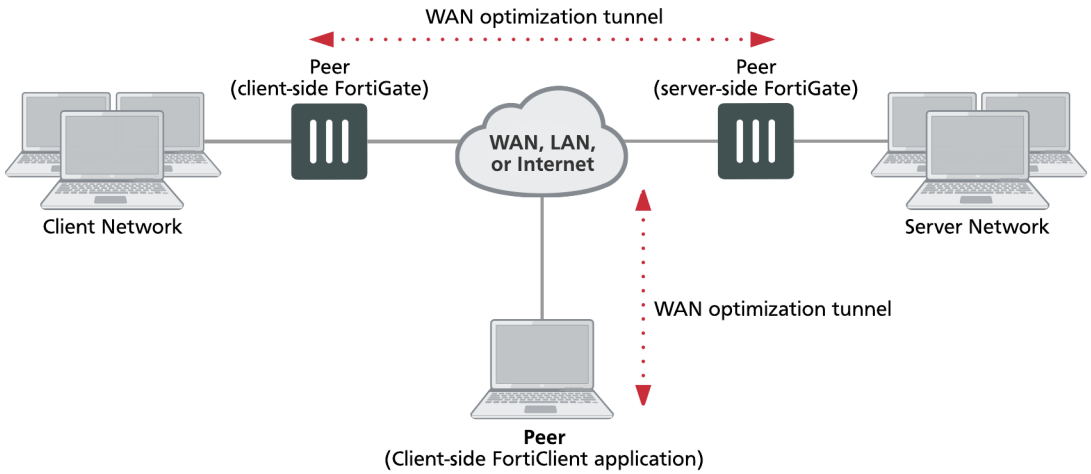

To identify all of the WAN optimization peers that a FortiGate unit can perform WAN optimization with, you add host IDs and IP addresses of all of the peers to the FortiGate unit configuration. The peer IP address is actually the IP address of the peer unit interface that communicates with the FortiGate unit.

Manual (peer-to-peer) and active-passive WAN optimization

You can create **manual** (peer-to-peer) and **active-passive** WAN optimization configurations.

In reality, because WAN optimization traffic can only be processed by one CPU core, it is not recommended to increase the number of manual mode peers on the FortiGate unit per VDOM.

Note that the maximum number of manual peers are restricted to 256 per VDOM. However, in Active-Passive configurations, there is no hard-limit to the maximum number of manual peers per VDOM.

Manual (peer to peer) configurations

Manual configurations allow for WAN optimization between one client-side FortiGate unit and one server-side FortiGate unit. To create a manual configuration you add a **manual mode** WAN optimization security policy to the client-side FortiGate unit. The manual mode policy includes the peer ID of a server-side FortiGate unit.

In a manual mode configuration, the client-side peer can only connect to the named server-side peer. When the client-side peer initiates a tunnel with the server-side peer, the packets that initiate the tunnel include extra information so that the server-side peer can determine that it is a peer-to-peer tunnel request. This extra information is required because the server-side peer does not require a WAN optimization policy; however, you need to add the client peer host ID and IP address to the server-side FortiGate unit peer list.

In addition, from the server-side FortiGate unit CLI you must and an Explicit Proxy security policy with `proxy` set to `wanopt` and the destination interface and network set to the network containing the servers that clients connect to over the WAN optimization tunnel. WAN optimization tunnel requests are accepted by the explicit proxy policy and if the client-side peer is in the server side peer's address list the traffic is forwarded to the servers on the destination network.

Manual mode client-side policy

You must configure manual mode client-side policies from the CLI. From the GUI a manual mode policy has WAN Optimization turned on and includes the following text beside the *WAN optimization* field: *Manual (Profile: <profile-name>. Peer: <peer-name>*.

Add a manual mode policy to the client-side FortiGate unit from the CLI. The policy enables WAN optimization, sets `wanopt-detection` to `off`, and uses the `wanopt-peer` option to specify the server-side peer. The following example uses the default WAN optimization profile.

```
config firewall policy
   edit 2
      set srcintf internal
      set dstintf wan1
      set srcaddr client-subnet
      set dstaddr server-subnet
      set action accept
      set schedule always
      set service ALL
      set wanopt enable
      set wanopt-detection off
      set wanopt-profile default
      set wanopt-peer server
   next
end
```

Manual mode server-side explicit proxy policy

The server-side explicit proxy policy allows connections from the WAN optimization tunnel to the server network by setting the proxy type to `wanopt`. You must add policies that set `proxy` to `wanopt` from the CLI and these policies do not appear on the GUI. The policy should look like the following:

```
configure firewall explicit-proxy-policy
   edit 3
      set proxy wanopt
      set dstintf internal
      set srcaddr all
      set dstaddr server-subnet
      set action accept
      set schedule always
      set service ALL
   next
end
```

Active-passive configurations

Active-passive WAN optimization requires an **active** WAN optimization policy on the client-side FortiGate unit and a **passive** WAN optimization policy on the server-side FortiGate unit. The server-side FortiGate unit also requires an explicit proxy policy with `proxy` set to `wanopt`.

You can use the passive policy to control WAN optimization address translation by specifying **transparent mode** or non-transparent mode. SeeWAN optimization transparent mode on page 35. You can also use the passive policy to apply security profiles, web caching, and other FortiGate features at the server-side FortiGate unit. For example, if a server-side FortiGate unit is protecting a web server, the passive policy could enable web caching.

A single passive policy can accept tunnel requests from multiple FortiGate units as long as the server-side FortiGate unit includes their peer IDs and all of the client-side FortiGate units include the server-side peer ID.

Active client-side policy

Add an active policy to the client-side FortiGate unit by turning on **WAN Optimization** and selecting **active**. Then select a WAN optimization **Profile.** From the CLI the policy could look like the following:

```
config firewall policy
   edit 2
      set srcintf internal
      set dstintf wan1
      set srcaddr client-subnet
      set dstaddr server-subnet
      set action accept
      set schedule always
      set service ALL
      set wanopt enable
      set wanopt-detection active
      set wanopt-profile default
   next
end
```

Server-side tunnel policy

The server-side requires an explicit proxy policy that sets the `proxy` to `wanopt`. You must add this policy from the CLI and policies with `proxy` set to `wanopt` do not appear on the GUI. From the CLI the policy could look like

the following:

```
configure firewall explicit-proxy-policy
   edit 3
      set proxy wanopt
      set dstintf internal
      set srcaddr all
      set dstaddr server-subnet
      set action accept
      set schedule always
      set service ALL
   next
end
```

Server-side passive policy

Add a passive policy to the server-side FortiGate unit by selecting **Enable WAN Optimization** and selecting **passive**. Then set the **Passive Option** to **transparent**. From the CLI the policy could look like the following:

```
config firewall policy
edit 2
      set srcintf "wan1"
      set dstintf "internal"
      set srcaddr "all"
      set dstaddr "all"
      set action accept
      set schedule "always"
      set service "ANY"
      set wanopt enable
      set wanopt-detection passive
      set wanopt-passive-opt transparent
   next
```

WAN optimization profiles

Use WAN optimization profiles to apply WAN optimization techniques to traffic to be optimized. In a WAN optimization profile you can select the protocols to be optimized and for each protocol you can enable SSL offloading (if supported), secure tunneling, byte caching and set the port or port range the protocol uses. You can also enable transparent mode and optionally select an authentication group. You can edit the default WAN optimization profile or create new ones.

To configure a WAN optimization profile go to **WAN Opt. & Cache > Profiles** and edit a profile or create a new one.

Configuring a WAN optimization profile

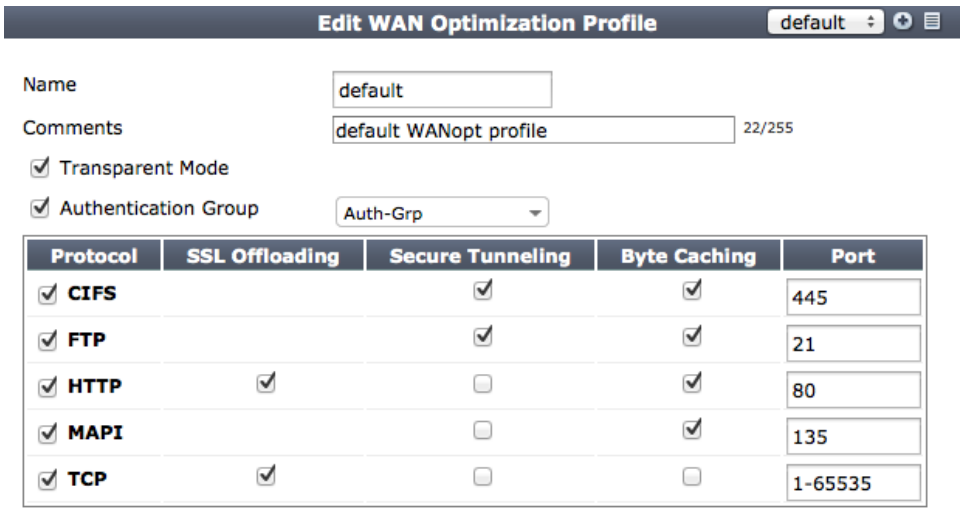

From the CLI you can use the following command to configure a WAN optimization profile to optimize HTTP traffic.

```
config wanopt profile
   edit new-profile
      config http
         set status enable
   end
```

Transparent Mode	Servers receiving packets after WAN optimization "see" different source addresses depending on whether or not you select **Transparent Mode**. For more information, see WAN optimization transparent mode on page 35.
Authentication Group	Select this option and select an authentication group so that the client and server-side FortiGate units must authenticate with each other before starting the WAN optimization tunnel. You must also select an authentication group if you select **Secure Tunneling** for any protocol. You must add identical authentication groups to both of the FortiGate units that will participate in the WAN optimization tunnel. For more information, see Configuring authentication groups on page 47.
Protocol	Select CIFS, FTP, HTTP or MAPI to apply protocol optimization for the selected protocols. See Protocol optimization on page 34. Select TCP if the WAN optimization tunnel accepts sessions that use more than one protocol or that do not use the CIFS, FTP, HTTP, or MAPI protocol.

SSL Offloading	Select to apply SSL offloading for HTTPS or other SSL traffic. You can use SSL offloading to offload SSL encryption and decryption from one or more HTTP servers to the FortiGate unit. If you enable this option, you must configure the security policy to accept SSL-encrypted traffic. If you enable SSL offloading, you must also use the CLI command `config firewall ssl-server` to add an SSL server for each HTTP server that you want to offload SSL encryption/decryption for. For more information, see Turning on web caching for HTTPS traffic on page 74.
Secure Tunnelling	The WAN optimization tunnel is encrypted using SSL encryption. You must also add an authentication group to the profile. For more information, see Secure tunneling on page 49.
Byte Caching	Select to apply WAN optimization byte caching to the sessions accepted by this rule. For more information, see "Byte caching".
Port	Enter a single port number or port number range. Only packets whose destination port number matches this port number or port number range will be optimized.

Processing non-HTTP sessions accepted by a WAN optimization profile with HTTP optimization

From the CLI, you can use the following command to configure how to process non-HTTP sessions when a rule configured to accept and optimize HTTP traffic accepts a non-HTTP session. This can occur if an application sends non-HTTP sessions using an HTTP destination port.

```
config wanopt profile
   edit default
      config http
         set status enable
         set tunnel-non-http {disable | enable}
      end
```

To drop non-HTTP sessions accepted by the rule set `tunnel-non-http` to `disable`, or set it to `enable` to pass non-HTTP sessions through the tunnel without applying protocol optimization, byte-caching, or web caching. In this case, the FortiGate unit applies TCP protocol optimization to non-HTTP sessions.

Processing unknown HTTP sessions

Unknown HTTP sessions are HTTP sessions that do not comply with HTTP 0.9, 1.0, or 1.1. From the CLI, use the following command to specify how a rule handles such HTTP sessions.

```
config wanopt profile
   edit default
      config http
         set status enable
         set unknown-http-version {best-effort | reject | tunnel}
      end
```

To assume that all HTTP sessions accepted by the rule comply with HTTP 0.9, 1.0, or 1.1, select `best-effort`. If a session uses a different HTTP version, WAN optimization may not parse it correctly. As a result, the FortiGate unit may stop forwarding the session and the connection may be lost. To reject HTTP sessions that do not use HTTP 0.9, 1.0, or 1.1, select `reject`.

To pass HTTP sessions that do not use HTTP 0.9, 1.0, or 1.1, but without applying HTTP protocol optimization, byte-caching, or web caching, you can also select `tunnel`. TCP protocol optimization is applied to these HTTP sessions.

Protocol optimization

Protocol optimization techniques optimize bandwidth use across the WAN. These techniques can improve the efficiency of communication across the WAN optimization tunnel by reducing the amount of traffic required by communication protocols. You can apply protocol optimization to Common Internet File System (CIFS), FTP, HTTP, MAPI, and general TCP sessions. You can apply general TCP optimization to MAPI sessions.

For example, CIFS provides file access, record locking, read/write privileges, change notification, server name resolution, request batching, and server authentication. CIFS is a fairly “chatty” protocol, requiring many background transactions to successfully transfer a single file. This is usually not a problem across a LAN. However, across a WAN, latency and bandwidth reduction can slow down CIFS performance.

When you select the CIFS protocol in a WAN optimization profile, the FortiGate units at both ends of the WAN optimization tunnel use a number of techniques to reduce the number of background transactions that occur over the WAN for CIFS traffic.

If a policy accepts a range of different types of traffic, you can set **Protocol** to **TCP** to apply general optimization techniques to TCP traffic. However, applying this TCP optimization is not as effective as applying more protocol-specific optimization to specific types of traffic. TCP protocol optimization uses techniques such as TCP SACK support, TCP window scaling and window size adjustment, and TCP connection pooling to remove TCP bottlenecks.

Protocol optimization and MAPI

By default the MAPI service uses port number 135 for RPC port mapping and may use random ports for MAPI messages. The random ports are negotiated through sessions using port 135. The FortiOS DCE-RPC session helper learns these ports and opens pinholes for the messages. WAN optimization is also aware of these ports and attempts to apply protocol optimization to MAPI messages that use them. However, to configure protocol optimization for MAPI you should set the WAN optimization profile to a single port number (usually port 135). Specifying a range of ports may reduce performance.

Byte caching

Byte caching breaks large units of application data (for example, a file being downloaded from a web page) into small chunks of data, labelling each chunk of data with a hash of the chunk and storing those chunks and their hashes in a database. The database is stored on a WAN optimization storage device. Then, instead of sending the actual data over the WAN tunnel, the FortiGate unit sends the hashes. The FortiGate unit at the other end of

the tunnel receives the hashes and compares them with the hashes in its local byte caching database. If any hashes match, that data does not have to be transmitted over the WAN optimization tunnel. The data for any hashes that does not match is transferred over the tunnel and added to that byte caching database. Then the unit of application data (the file being downloaded) is reassembled and sent to its destination.

The stored byte caches are not application specific. Byte caches from a file in an email can be used to optimize downloading that same file or a similar file from a web page.

The result is less data transmitted over the WAN. Initially, byte caching may reduce performance until a large enough byte caching database is built up.

To enable byte caching, you select **Byte Caching** in a WAN optimization profile.

Byte caching cannot determine whether or not a file is compressed (for example a zip file), and caches compressed and non-compressed versions of the same file separately.

Dynamic data chunking for byte caching

Dynamic data chunking can improve byte caching by improving detection of data chunks that are already cached in changed files or in data embedded in traffic using an unknown protocol. Dynamic data chunking is available for HTTP, CIFS and FTP.

Use the following command to enable dynamic data chunking for HTTP in the default WAN optimization profile.

```
config wanopt profile
   edit default
        config http
           set prefer-chunking dynamic
      end
```

By default dynamic data chunking is disabled and `prefer-chunking` is set to `fix`.

WAN optimization transparent mode

WAN optimization is transparent to users. This means that with WAN optimization in place, clients connect to servers in the same way as they would without WAN optimization. However, servers receiving packets after WAN optimization “see” different source addresses depending on whether or not transparent mode is selected for WAN optimization. If transparent mode is selected, WAN optimization keeps the original source address of the packets, so servers appear to receive traffic directly from clients. Routing on the server network should be configured to route traffic with client source IP addresses from the server-side FortiGate unit to the server and back to the server-side FortiGate unit.

Some protocols, for example CIFS, may not function as expected if transparent mode is **not** selected. In most cases, for CIFS WAN optimization you should select transparent mode and make sure the server network can route traffic as described to support transparent mode.

If transparent mode is not selected, the source address of the packets received by servers is changed to the address of the server-side FortiGate unit interface that sends the packets to the servers. So servers appear to receive packets from the server-side FortiGate unit. Routing on the server network is simpler in this case because

client addresses are not involved. All traffic appears to come from the server-side FortiGate unit and not from individual clients.

Do not confuse WAN optimization transparent mode with FortiGate transparent mode. WAN optimization transparent mode is similar to source NAT. FortiGate Transparent mode is a system setting that controls how the FortiGate unit (or a VDOM) processes traffic.

Configuring Transparent mode

You can configure transparent mode by selecting **Transparent** in a WAN Optimization profile. The profile is added to an active WAN Optimization policy.

When you configure a passive WAN Optmization policy you can accept the active policy transparent setting or you can override the active policy transparent setting. From the GUI you can do this by setting the **Passive Option** as follows:

- **default** use the transparent setting in the WAN Optimization profile added to the active policy (client-side configuration).
- **transparent** impose transparent mode (override the active policy transparent mode setting). Packets exiting the FortiGate keep their original source addresses.
- **non-transparent** impose non-transparent mode (override the active policy transparent mode setting). Packets exiting the FortiGate have their source address changed to the address of the server-side FortiGate unit interface that sends the packets to the servers.

From the CLI you can use the following command:

```
config firewall policy
   set wanopt-passive-opt {default | transparent | non-transparent}
end
```

FortiClient WAN optimization

PCs running the FortiClient application are client-side peers that initiate WAN optimization tunnels with server-side peer FortiGate units. However, you can have an ever-changing number of FortiClient peers with IP addresses that also change regularly. To avoid maintaining a list of such peers, you can instead configure WAN optimization to accept any peer and use authentication to identify FortiClient peers.

Together, the WAN optimization peers apply the WAN optimization features to optimize the traffic flow over the WAN between the clients and servers. WAN optimization reduces bandwidth requirements, increases throughput, reduces latency, offloads SSL encryption/decryption and improves privacy for traffic on the WAN.

For more details, see FortiClient WAN optimization on page 91.

Operating modes and VDOMs

To use WAN optimization, the FortiGate units can operate in either NAT/Route or Transparent mode. The client-side and server-side FortiGate units do not have to be operating in the same mode.

As well, the FortiGate units can be configured for multiple virtual domain (VDOM) operation. You configure WAN optimization for each VDOM and configure one or both of the units to operate with multiple VDOMs enabled.

If a FortiGate unit or VDOM is operating in Transparent mode with WAN optimization enabled, WAN optimization uses the management IP address as the peer IP address of the FortiGate unit instead of the address of an interface.

WAN optimization tunnels

All optimized traffic passes between the FortiGate units or between a FortiClient peer and a FortiGate unit over a WAN optimization tunnel. Traffic in the tunnel can be sent in plain text or encrypted using AES-128bit-CBC SSL.

WAN optimization tunnels

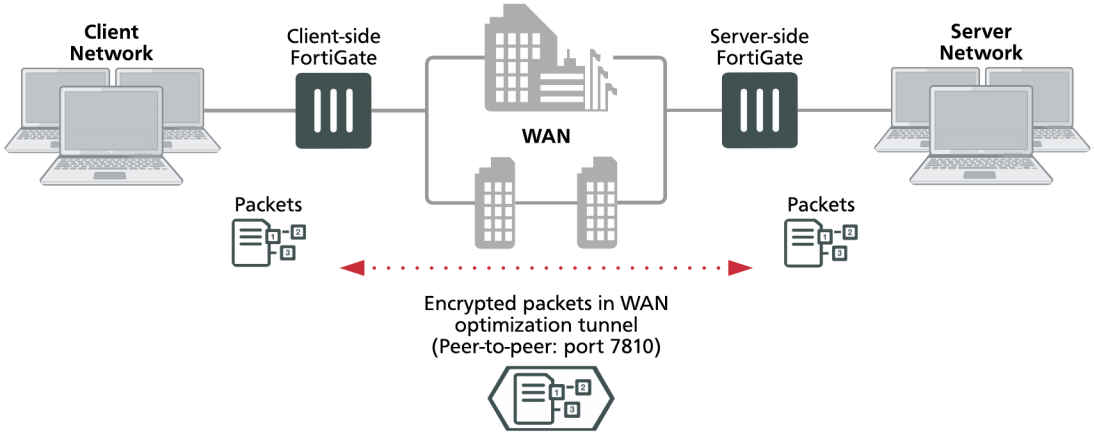

Both plain text and the encrypted tunnels use TCP destination port 7810.

Before a tunnel can be started, the peers must be configured to authenticate with each other. Then, the client-side peer attempts to start a WAN optimization tunnel with the server-side peer. Once the peers authenticate with each other, they bring up the tunnel and WAN optimization communication over the tunnel starts. After a tunnel has been established, multiple WAN optimization sessions can start and stop between peers without restarting the tunnel.

Tunnel sharing

You can use the `tunnel-sharing` WAN optimization profile CLI keyword to configure tunnel sharing for WAN optimization rules. Tunnel sharing means multiple WAN optimization sessions share the same tunnel. Tunnel sharing can improve performance by reducing the number of WAN optimization tunnels between FortiGate units. Having fewer tunnels means less data to manage. Also, tunnel setup requires more than one exchange of information between the ends of the tunnel. Once the tunnel is set up, each new session that shares the tunnel avoids tunnel setup delays.

Tunnel sharing also uses bandwidth more efficiently by reducing the chances that small packets will be sent down the tunnel. Processing small packets reduces network throughput, so reducing the number of small packets improves performance. A shared tunnel can combine all the data from the sessions being processed by the tunnel and send the data together. For example, suppose a FortiGate unit is processing five WAN optimization sessions and each session has 100 bytes to send. If these sessions use a shared tunnel, WAN optimization combines the

packets from all five sessions into one 500-byte packet. If each session uses its own private tunnel, five 100-byte packets will be sent instead. Each packet also requires a TCP ACK reply. The combined packet in the shared tunnel requires one TCP ACK packet. The separate packets in the private tunnels require five.

Use the following command to configure tunnel sharing for HTTP traffic in a WAN optimization profile.

```
config wanopt profile
   edit default
      config http
         set tunnel-sharing {express-shared | private | shared}
      end
```

Tunnel sharing is not always recommended and may not always be the best practice. Aggressive and non-aggressive protocols should not share the same tunnel. An aggressive protocol can be defined as a protocol that is able to get more bandwidth than a non-aggressive protocol. (The aggressive protocols can “starve” the non-aggressive protocols.) HTTP and FTP are considered aggressive protocols. If aggressive and non-aggressive protocols share the same tunnel, the aggressive protocols may take all of the available bandwidth. As a result, the performance of less aggressive protocols could be reduced. To avoid this problem, rules for HTTP and FTP traffic should have their own tunnel. To do this, set `tunnel-sharing` to `private` for WAN optimization rules that accept HTTP or FTP traffic.

It is also useful to set `tunnel-sharing` to `express-shared` for applications, such as Telnet, that are very interactive but not aggressive. Express sharing optimizes tunnel sharing for Telnet and other interactive applications where latency or delays would seriously affect the user’s experience with the protocol.

Set `tunnel-sharing` to `shared` for applications that are not aggressive and are not sensitive to latency or delays. WAN optimization rules set to `sharing` and `express-shared` can share the same tunnel.

WAN optimization and user and device identity policies, load balancing and traffic shaping

Please note the following about WAN optimization and firewall policies:

- WAN optimization is not compatible with firewall load balancing.
- WAN optimization is compatible with source and destination NAT options in firewall policies (including firewall virtual IPs). If a virtual IP is added to a policy the traffic that exits the WAN optimization tunnel has its destination address changed to the virtual IPs mapped to IP address and port.
- WAN optimization is compatible with user identity-based and device identity security policies. If a session is allowed after authentication or device identification the session can be optimized.

Traffic shaping

Traffic shaping works for WAN optimization traffic that is not in a WAN optimization tunnel. So traffic accepted by a WAN optimization security policy on a client-side FortiGate unit can be shaped on ingress. However, when the traffic enters the WAN optimization tunnel, traffic shaping is not applied.

In manual mode:

- Traffic shaping works as expected on the client-side FortiGate unit.
- Traffic shaping cannot be applied to traffic on the server-side FortiGate unit.

In active-passive mode:

- Traffic shaping works as expected on the client-side FortiGate unit.
- If transparent mode is enabled in the WAN optimization profile, traffic shaping also works as expected on the server-side FortiGate unit.
- If transparent mode is not enabled, traffic shaping works partially on the server-side FortiGate unit.

WAN optimization and HA

You can configure WAN optimization on a FortiGate HA cluster. The recommended best practice HA configuration for WAN optimization is active-passive mode. When the cluster is operating, all WAN optimization sessions are processed by the primary unit only. Even if the cluster is operating in active-active mode, HA does not load-balance WAN optimization sessions.

You can also form a WAN optimization tunnel between a cluster and a standalone FortiGate unit or between two clusters.

In a cluster, only the primary unit stores the byte cache database. This database is not synchronized to the subordinate units. So, after a failover, the new primary unit must rebuild its byte cache. Rebuilding the byte cache can happen relatively quickly because the new primary unit gets byte cache data from the other FortiGate unit that it is participating with in WAN optimization tunnels.

WAN optimization, web caching and memory usage

To accelerate and optimize disk access and to provide better throughput and less latency FortiOS WAN optimization uses provisioned memory to reduce disk I/O and increase disk I/O efficiency. In addition, WAN optimization requires a small amount of additional memory per session for comprehensive flow control logic and efficient traffic forwarding.

When WAN optimization is enabled you will see a reduction in available memory. The reduction increases when more WAN optimization sessions are being processed. If you are thinking of enabling WAN optimization on an operating FortiGate unit, make sure its memory usage is not maxed out during high traffic periods.

In addition to using the system dashboard to see the current memory usage you can use the `get test wad 2` command to see how much memory is currently being used by WAN optimization. See "get test {wad | wccpd} <test_level>" for more information.

Monitoring WAN optimization performance

Using WAN optimization monitoring, you can confirm that a FortiGate unit is optimizing traffic and view estimates of the amount of bandwidth saved. The WAN optimization monitor presents collected log information in a graphical format to show network traffic summary and bandwidth optimization information.

To view the WAN optimization monitor, go to **Monitor > WAN Opt. Monitor**.

WAN optimization monitor

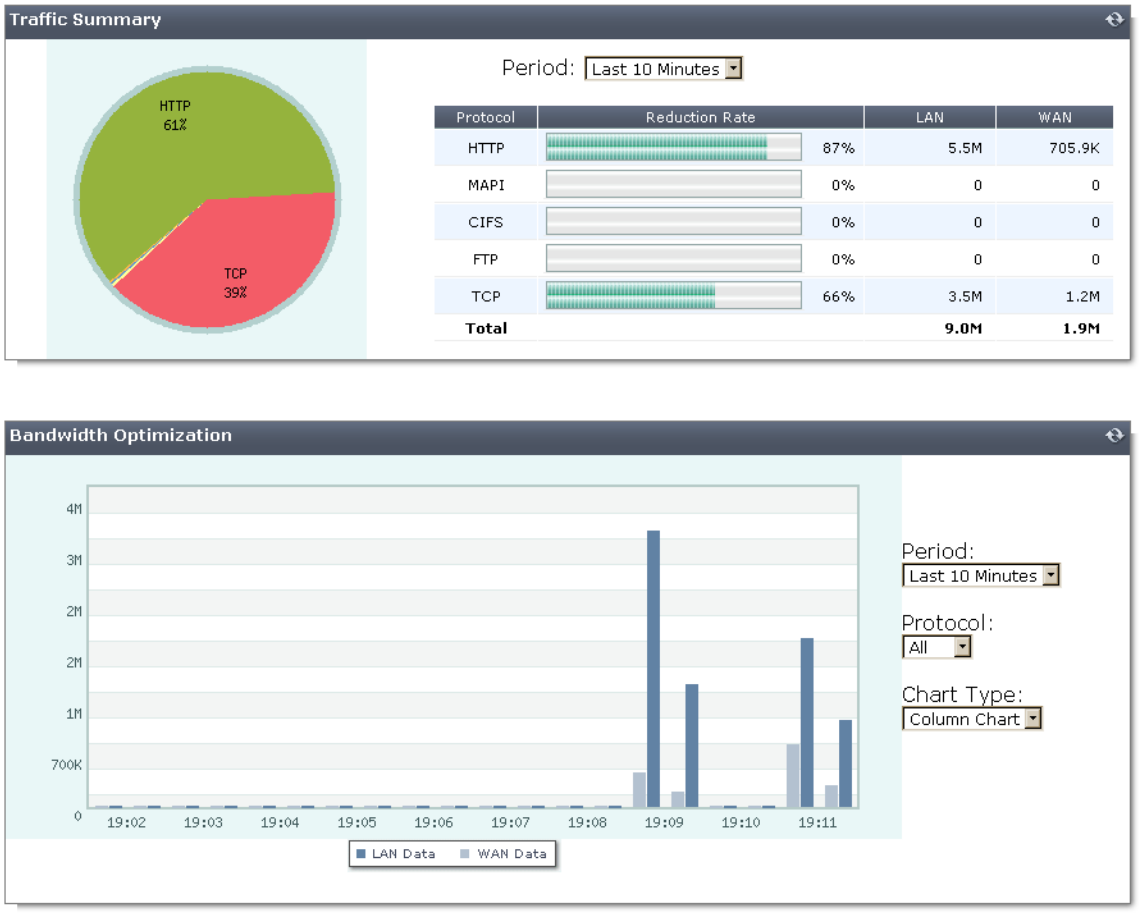

Traffic Summary

The traffic summary shows how WAN optimization is reducing the amount of traffic on the WAN for each WAN optimization protocol by showing the traffic reduction rate as a percentage of the total traffic. The traffic summary also shows the amount of WAN and LAN traffic. If WAN optimization is being effective the amount of WAN traffic should be lower than the amount of LAN traffic.

You can use the refresh icon to update the traffic summary display at any time. You can also set the amount of time for which the traffic summary shows data. The time period can vary from the last 10 minutes to the last month.

Bandwidth Optimization

This section shows network bandwidth optimization per time period. A line or column chart compares an application's pre-optimized size (LAN data) with its optimized size (WAN data). You can select the chart type, the monitoring time period, and the protocol for which to display data. If WAN optimization is being effective the WAN bandwidth should be lower than the LAN bandwidth.

WAN optimization configuration summary

This section includes a client-side and a server-side WAN Optimization configuration summary.:

Client-side configuration summary

WAN optimization profile

Enter the following command to view WAN optimization profile CLI options:

```
tree wanopt profile
-- [profile] --*name (36)
      |- transparent
      |- comments
      |- auth-group (36)
      |- <http> -- status
            |- secure-tunnel
            |- byte-caching
            |- prefer-chunking
            |- tunnel-sharing
            |- log-traffic
            |- port (1,65535)
            |- ssl
            |- ssl-port (1,65535)
            |- unknown-http-version
            +- tunnel-non-http
      |- <cifs> -- status
            |- secure-tunnel
            |- byte-caching
            |- prefer-chunking
            |- tunnel-sharing
            |- log-traffic
            +- port (1,65535)
      |- <mapi> -- status
            |- secure-tunnel
            |- byte-caching
            |- tunnel-sharing
            |- log-traffic
            +- port (1,65535)
      |- <ftp> -- status
            |- secure-tunnel
            |- byte-caching
            |- prefer-chunking
            |- tunnel-sharing
            |- log-traffic
            +- port (1,65535)
      +- <tcp> -- status
            |- secure-tunnel
            |- byte-caching
            |- byte-caching-opt
            |- tunnel-sharing
            |- log-traffic
            |- port
            |- ssl
```

```
            +- ssl-port (1,65535)
```

Local host ID and peer settings

```
config wanopt settings
   set host-id client
end
config wanopt peer
   edit server
   set ip 10.10.2.82
end
```

Security policies

Two client-side WAN optimization security policy configurations are possible. One for active-passive WAN optimization and one for manual WAN optimization.

Active/passive mode on the client-side

```
config firewall policy
   edit 2
      set srcintf internal
      set dstintf wan1
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service ALL
      set wanopt enable <<< enable WAN optimization
      set wanopt-detection active <<< set the mode to active/passive
      set wanopt-profile "default" <<< select the wanopt profile
   next
end
```

Manual mode on the client-side

```
config firewall policy
   edit 2
      set srcintf internal
      set dstintf wan1
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service ALL
      set wanopt enable <<< enable WAN optimization
      set wanopt-detection off <<< sets the mode to manual
      set wanopt-profile "default" <<< select the wanopt profile
      set wanopt-peer "server" <<< set the only peer to do wanopt
                                                                    with
          (required for manual mode)
   next
end
```

server-side configuration summary

Local host ID and peer settings

```
config wanopt settings
   set host-id server
end
config wanopt peer
   edit client
      set ip 10.10.2.81
end
```

Security policies

Two server-side WAN optimization security policy configurations are possible. One for active-passive WAN optimization and one for manual WAN optimization.

Active/passive mode on server-side

```
config firewall policy
edit 2 <<< the passive mode policy
      set srcintf wan1
      set dstintf internal
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service ALL
      set wanopt enable
      set wanopt-detection passive
      set wanopt-passive-opt transparent
   end
config firewall explicit-proxy-policy
   edit 3 <<< policy that accepts wanopt tunnel connections from the       server
      set proxy wanopt <<< wanopt proxy type
      set dstintf internal
      set srcaddr all
      set dstaddr server-subnet
      set action accept
      set schedule always
      set service ALL
   next
end
```

Manual mode on server-side

```
config firewall explicit-proxy-policy
   edit 3 <<< policy that accepts wanopt tunnel connections from the       client
      set proxy wanopt <<< wanopt proxy type
      set dstintf internal
      set srcaddr all
      set dstaddr server-subnet
      set action accept
      set schedule always
      set service ALL
```

```
    next
end
```

Best practices

This is a short list of WAN optimization and explicit proxy best practices.

- WAN optimization tunnel sharing is recommended for similar types of WAN optimization traffic. However, tunnel sharing for different types of traffic is not recommended. For example, aggressive and non-aggressive protocols should not share the same tunnel. See Tunnel sharing on page 37.
- Active-passive HA is the recommended HA configuration for WAN optimization. See WAN optimization and HA on page 39.
- Configure WAN optimization authentication with specific peers. Accepting any peer is not recommended as this can be less secure. SeeAccepting any peers on page 45.
- Set the explicit proxy **Default Firewall Policy Action** to **Deny**. This means that a security policy is required to use the explicit web proxy. See The FortiGate explicit web proxy on page 94.
- Set the explicit FTP proxy **Default Firewall Policy Action** to **Deny**. This means that a security policy is required to use the explicit FTP proxy. See General explicit FTP proxy configuration steps on page 123.
- Do not enable the explicit web or FTP proxy on an interface connected to the Internet. This is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address. If you must enable the proxy on such an interface make sure authentication is required to use the proxy. See General explicit web proxy configuration steps on page 95.

Peers and authentication groups

All communication between WAN optimization peers begins with one WAN optimization peer (or client-side FortiGate unit) sending a WAN optimization tunnel request to another peer (or server-side FortiGate unit). During this process, the WAN optimization peers identify and optionally authenticate each other.

Basic WAN optimization peer requirements

WAN optimization requires the following configuration on each peer. For information about configuring local and peer host IDs, see Configuring peers on page 46.

- The peer must have a unique host ID.
- Unless authentication groups are used, peers authenticate each other using host ID values. Do not leave the local host ID at its default value.
- The peer must know the host IDs and IP addresses of all of the other peers that it can start WAN optimization tunnels with. This does not apply if you use authentication groups that accept all peers.
- All peers must have the same local certificate installed on their FortiGate units if the units authenticate by local certificate. Similarly, if the units authenticate by pre-shared key (password), administrators must know the password. The type of authentication is selected in the authentication group. This applies only if you use authentication groups.

Accepting any peers

Strictly speaking, you do not need to add peers. Instead you can configure authentication groups that accept any peer. However, for this to work, both peers must have the same authentication group (with the same name) and both peers must have the same certificate or pre-shared key.

Accepting any peer is useful if you have many peers or if peer IP addresses change. For example, you could have many travelling FortiClient peers with IP addresses that are always changing as the users travel to different customer sites. This configuration is also useful if you have FortiGate units with dynamic external IP addresses (using DHCP or PPPoE). For most other situations, this method is not recommended and is not a best practice as it is less secure than accepting defined peers or a single peer. For more information, see Configuring authentication groups on page 47.

How FortiGate units process tunnel requests for peer authentication

When a client-side FortiGate unit attempts to start a WAN optimization tunnel with a peer server-side FortiGate unit, the tunnel request includes the following information:

- the client-side local host ID
- the name of an authentication group, if included in the rule that initiates the tunnel
- if an authentication group is used, the authentication method it specifies: pre-shared key or certificate
- the type of tunnel (secure or not).

For information about configuring the local host ID, peers and authentication groups, see Configuring peers on page 46 and Configuring authentication groups on page 47.

The authentication group is optional unless the tunnel is a secure tunnel. For more information, see Secure tunneling on page 49.

If the tunnel request includes an authentication group, the authentication will be based on the settings of this group as follows:

- The server-side FortiGate unit searches its own configuration for the name of the authentication group in the tunnel request. If no match is found, the authentication fails.
- If a match is found, the server-side FortiGate unit compares the authentication method in the client and server authentication groups. If the methods do not match, the authentication fails.
- If the authentication methods match, the server-side FortiGate unit tests the peer acceptance settings in its copy of the authentication group.
- If the setting is **Accept Any Peer**, the authentication is successful.
- If the setting is **Specify Peer**, the server-side FortiGate unit compares the client-side local host ID in the tunnel request with the peer name in the server-side authentication group. If the names match, authentication is successful. If a match is not found, authentication fails.
- If the setting is **Accept Defined Peers**, the server-side FortiGate unit compares the client-side local host ID in the tunnel request with the server-side peer list. If a match is found, authentication is successful. If a match is not found, authentication fails.

If the tunnel request does not include an authentication group, authentication will be based on the client-side local host ID in the tunnel request. The server-side FortiGate unit searches its peer list to match the client-side local host ID in the tunnel request. If a match is found, authentication is successful. If a match is not found, authentication fails.

If the server-side FortiGate unit successfully authenticates the tunnel request, the server-side FortiGate unit sends back a tunnel setup response message. This message includes the server-side local host ID and the authentication group that matches the one in the tunnel request.

The client-side FortiGate unit then performs the same authentication procedure as the server-side FortiGate unit did. If both sides succeed, tunnel setup continues.

Configuring peers

When you configure peers, you first need to add the local host ID that identifies the FortiGate unit for WAN optimization and then add the peer host ID and IP address of each FortiGate unit with which a FortiGate unit can create WAN optimization tunnels.

To configure WAN optimization peers - web-based manager:

1. Go to **WAN Opt. & Cache > Peers**.
2. For **Local Host ID**, enter the local host ID of **this** FortiGate unit and select **Apply**. If you add this FortiGate unit as a peer to another FortiGate unit, use this ID as its **peer** host ID.

   The local or host ID can contain up to 25 characters and can include spaces.
3. Select **Create New** to add a new peer.

4. For **Peer Host ID**, enter the peer host ID of the peer FortiGate unit. This is the local host ID added to the peer FortiGate unit.
5. For **IP Address**, add the IP address of the peer FortiGate unit. This is the source IP address of tunnel requests sent by the peer, usually the IP address of the FortiGate interface connected to the WAN.
6. Select **OK**.

To configure WAN optimization peers - CLI:

In this example, the local host ID is named `HQ_Peer` and has an IP address of `172.20.120.100`. Three peers are added, but you can add any number of peers that are on the WAN.

1. Enter the following command to set the local host ID to `HQ_Peer`.

```
config wanopt settings
   set host-id HQ_peer
end
```

2. Enter the following commands to add three peers.

```
config wanopt peer
   edit Wan_opt_peer_1
      set ip 172.20.120.100
   next
   edit Wan_opt_peer_2
      set ip 172.30.120.100
   next
   edit Wan_opt_peer_3
      set ip 172.40.120.100
   end
```

Configuring authentication groups

You need to add authentication groups to support authentication and secure tunneling between WAN optimization peers.

To perform authentication, WAN optimization peers use a certificate or a pre-shared key added to an authentication group so they can identify each other before forming a WAN optimization tunnel. Both peers must have an authentication group with the same name and settings. You add the authentication group to a peer-to-peer or active rule on the client-side FortiGate unit. When the server-side FortiGate unit receives a tunnel start request from the client-side FortiGate unit that includes an authentication group, the server-side FortiGate unit finds an authentication group in its configuration with the same name. If both authentication groups have the same certificate or pre-shared key, the peers can authenticate and set up the tunnel.

Authentication groups are also required for secure tunneling.

To add authentication groups, go to **WAN Opt. & Cache > Authentication Groups**.

To add an authentication group - web-based manager:

Use the following steps to add any kind of authentication group. It is assumed that if you are using a local certificate to authenticate, it is already added to the FortiGate unit

1. Go to **WAN Opt. & Cache > Authentication Groups**.
2. Select **Create New**.

3. Add a **Name** for the authentication group.

   You will select this name when you add the authentication group to a WAN optimization rule.
4. Select the **Authentication Method**.

   Select **Certificate** if you want to use a certificate to authenticate and encrypt WAN optimization tunnels. You must select a local certificate that has been added to this FortiGate unit. (To add a local certificate, go to **System > Certificates**.) Other FortiGate units that participate in WAN optimization tunnels with this FortiGate unit must have an authentication group with the same name and certificate.

   Select **Pre-shared key** if you want to use a pre-shared key or password to authenticate and encrypt WAN optimization tunnels. You must add the **Password** (or pre-shared key) used by the authentication group. Other FortiGate units that participate in WAN optimization tunnels with this FortiGate unit must have an authentication group with the same name and password. The password must contain at least 6 printable characters and should be known only by network administrators. For optimum protection against currently known attacks, the key should consist of a minimum of 16 randomly chosen alphanumeric characters.
5. Configure **Peer Acceptance** for the authentication group.

   Select **Accept Any Peer** if you do not know the peer host IDs or IP addresses of the peers that will use this authentication group. This setting is most often used for WAN optimization with the FortiClient application or with FortiGate units that do not have static IP addresses, for example units that use DHCP.

   Select **Accept Defined Peers** if you want to authenticate with peers added to the peer list only.

   Select **Specify Peer** and select one of the peers added to the peer list to authenticate with the selected peer only.
6. Select **OK**.
7. Add the authentication group to a WAN optimization rule to apply the authentication settings in the authentication group to the rule.

To add an authentication group that uses a certificate- CLI:

Enter the following command to add an authentication group that uses a certificate and can authenticate all peers added to the FortiGate unit configuration.

In this example, the authentication group is named `auth_grp_1` and uses a certificate named `Example_Cert`.

```
config wanopt auth-group
   edit auth_grp_1
      set auth-method cert
      set cert Example_Cert
      set peer-accept defined
   end
```

To add an authentication group that uses a pre-shared key - CLI:

Enter the following command to add an authentication group that uses a pre-shared key and can authenticate only the peer added to the authentication group.

In this example, the authentication group is named `auth_peer`, the peer that the group can authenticate is named `Server_net`, and the authentication group uses `123456` as the pre-shared key. In practice you should use a more secure pre-shared key.

```
config wanopt auth-group
   edit auth_peer
      set auth-method psk
      set psk 123456
      set peer-accept one
      set peer Server_net
   end
```

To add an authentication group that accepts WAN optimization connections from any peer - web-based manager

Add an authentication group that accepts any peer for situations where you do not have the **Peer Host IDs** or **IP Addresses** of the peers that you want to perform WAN optimization with. This setting is most often used for WAN optimization with the FortiClient application or with FortiGate units that do not have static IP addresses, for example units that use DHCP. An authentication group that accepts any peer is less secure than an authentication group that accepts defined peers or a single peer.

The example below sets the authentication method to **Pre-shared key**. You must add the same password to all FortiGate units using this authentication group.

1. Go to **WAN Opt. & Cache > Authentication Groups**.
2. Select **Create New** to add a new authentication group.
3. Configure the authentication group:

Name	Specify any name.
Authentication Method	Pre-shared key
Password	Enter a pre-shared key.
Peer Acceptance	Accept Any Peer

To add an authentication group that accepts WAN optimization connections from any peer - CLI:

In this example, the authentication group is named `auth_grp_1`. It uses a certificate named `WAN_Cert` and accepts any peer.

```
config wanopt auth-group
   edit auth_grp_1
      set auth-method cert
      set cert WAN_Cert
      set peer-accept any
   end
```

Secure tunneling

You can configure WAN optimization rules to use AES-128bit-CBC SSL to encrypt the traffic in the WAN optimization tunnel. WAN optimization uses FortiASIC acceleration to accelerate SSL decryption and encryption

of the secure tunnel. Peer-to-peer secure tunnels use the same TCP port as non-secure peer-to-peer tunnels (TCP port 7810).

To use secure tunneling, you must select **Enable Secure Tunnel** in a WAN optimization rule and add an authentication group. The authentication group specifies the certificate or pre-shared key used to set up the secure tunnel. The **Peer Acceptance** setting of the authentication group does not affect secure tunneling.

The FortiGate units at each end of the secure tunnel must have the same authentication group with the same name and the same configuration, including the same pre-shared key or certificate. To use certificates you must install the same certificate on both FortiGate units.

For active-passive WAN optimization you can select **Enable Secure Tunnel** only in the active rule. In peer-to-peer WAN optimization you select **Enable Secure Tunnel** in the WAN optimization rule on both FortiGate units. For information about active-passive and peer-to-peer WAN optimization, see Manual (peer-to-peer) and active-passive WAN optimization on page 28

For a secure tunneling configuration example, see Example Adding secure tunneling to an active-passive WAN optimization configuration on page 65.

Monitoring WAN optimization peer performance

The WAN optimization peer monitor lists all of the WAN optimization peers that a FortiGate unit can perform WAN optimization with. These include peers manually added to the configuration as well as discovered peers.

The monitor lists each peer's name, IP address, and peer type. The peer type indicates whether the peer was manually added or discovered. To show WAN optimization performance, for each peer the monitor lists the percent of traffic reduced by the peer in client-side WAN optimization configurations and in server-side configurations (also called gateway configurations).

To view the peer monitor, go to **WAN Opt. & Cache > Peer Monitor**.

Configuration examples

This chapter provides the basic examples to illustrate WAN optimization configurations introduced in the previous chapters.

Example Basic manual (peer-to-peer) WAN optimization configuration

In a manual (peer to peer) configuration the WAN optimization tunnel can be set up between one client-side FortiGate unit and one server-side FortiGate unit. The peer ID of the server-side FortiGate unit is added to the client-side WAN optimization policy. When the client-side FortiGate unit initiates a tunnel with the server-side FortiGate unit, the packets that initiate the tunnel include information that allows the server-side FortiGate unit to determine that it is a manual tunnel request. The server-side FortiGate unit does not require a WAN optimization profile; you just need to add the client peer host ID and IP address to the server-side FortiGate unit peer list and from the CLI an explicit proxy policy to accept WAN optimization tunnel connections.

In a manual WAN optimization configuration, you create a manual WAN optimization security policy on the client-side FortiGate unit. To do this you must use the CLI to set `wanopt-detection` to `off` and to add the peer host ID of the server-side FortiGate unit to the WAN optimization security policy.

Network topology and assumptions

This example configuration includes a client-side FortiGate unit called Client-Fgt with a WAN IP address of 172.20.34.12. This unit is in front of a network with IP address 172.20.120.0. The server-side FortiGate unit is called Server_Fgt with a WAN IP address of 192.168.30.12. This unit is in front of a web server network with IP address 192.168.10.0.

This example customizes the default WAN optimization profile on the client-side FortiGate unit and adds it to the WAN optimization policy. You can also create a new WAN optimization profile.

Example manual (peer-to-peer) topology

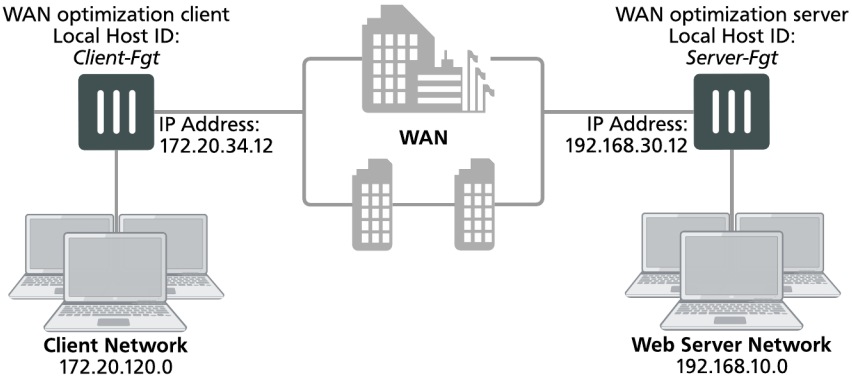

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Configure the client-side FortiGate unit:
   - Add peers.
   - Configure the default WAN optimization profile to optimize HTTP traffic.
   - Add a manual WAN optimization security policy.
2. Configure the server-side FortiGate unit:
   - Add peers.
   - Add a WAN optimization tunnel policy.

Configuring basic peer-to-peer WAN optimization - web-based manager

Use the following steps to configure the example configuration from the web-based manager.

To configure the client-side FortiGate unit

1. Go to **WAN Opt. & Cache > Peers** and enter a **Local Host ID** for the client-side FortiGate unit:

Local Host ID	Client-Fgt

2. Select **Apply**.
3. Select **Create New** and add the server-side FortiGate unit **Peer Host ID** and **IP Address** for the server-side FortiGate:

Peer Host ID	Server-Fgt
IP Address	192.168.30.12

4. Select **OK**.
5. Go to **Policy & Objects > Addresses** and select **Create New** to add a firewall address for the client network.

Category	Address
Name	Client-Net
Type	Subnet
Subnet / IP Range	172.20.120.0/24
Interface	port1

6. Select **Create New** to add a firewall address for the web server network.

Category	Address

Name	Web-Server-Net
Type	Subnet
Subnet / IP Range	192.168.10.0/24
Interface	port2

7. Go to **WAN Opt. & Cache > Profiles** and edit the default profile.
8. Select **Transparent Mode**.
9. Under Protocol, select **HTTP** and for HTTP select **Byte Caching**. Leave the HTTP **Port** set to 80.
10. Select **Apply** to save your changes.
11. Go to **Policy & Objects > IPv4 Policy** and add a WAN optimization security policy to the client-side FortiGate unit that accepts traffic to be optimized:

Incoming Interface	port1
Source Address	all
Outgoing Interface	port2
Destination Address	all
Schedule	always
Service	ALL
Action	ACCEPT

12. Select **Enable WAN Optimization** and configure the following settings:

Enable WAN Optimization	active
Profile	default

13. Select **OK**.
14. Edit the policy from the CLI to turn off `wanopt-detection`, add the peer ID of the server-side FortiGate unit, and the default WAN optimization profile. The following example assumes the ID of the policy is 5:

```
config firewall policy
  edit 5
    set wanopt-detection off
    set wanopt-peer Server-Fgt
    set wanopt-profile default
  end
```

When you set the detection mode to `off` the policy becomes a manual mode WAN optimization policy. On the web-based manager the WAN optimization part of the policy changes to the following:

Enable WAN Optimization	Manual (Profile: default, Peer: Peer-Fgt-2)

To configure the server-side FortiGate unit

1. Go to **WAN Opt. & Cache > Peers** and enter a **Local Host ID** for the server-side FortiGate unit:

Local Host ID	Server-Fgt

2. Select **Apply**.
3. Select **Create New** and add a **Peer Host ID** and the **IP Address** for the client-side FortiGate unit:

Peer Host ID	Client-Fgt
IP Address	172.20.34.12

4. Select **OK**.
5. Enter the following CLI command to add an explicit proxy policy to accept WAN optimization tunnel connections.

```
configure firewall explicit-proxy-policy
   edit 0
      set proxy wanopt
      set dstintf port1
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service ALL
   next
end
```

Configuring basic peer-to-peer WAN optimization - CLI

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit CLI.

To configure the client-side FortiGate unit

1. Add the Local Host ID to the client-side FortiGate configuration:

```
config wanopt settings
   set host-id Client-Fgt
end
```

2. Add the server-side Local Host ID to the client-side peer list:

```
config wanopt peer
   edit Server-Fgt
   set ip 192.168.30.12
end
```

3. Add a firewall address for the client network.

```
config firewall address
   edit Client-Net
      set type ipmask
      set subnet 172.20.120.0 255.255.255.0
      set associated-interface port1
   end
```

4. Add a firewall address for the web server network.

```
config firewall address
   edit Web-Server-Net
      set type ipmask
      set subnet 192.168.10.0 255.255.255.0
      set associated-interface port2
   end
```

5. Edit the default WAN optimization profile, select transparent mode, enable HTTP WAN optimization and enable byte caching for HTTP. Leave the HTTP Port set to 80.

```
config wanopt profile
   edit default
      set transparent enable
         config http
            set status enable
            set byte-caching enable
         end
   end
```

6. Add a WAN optimization security policy to the client-side FortiGate unit to accept the traffic to be optimized:

```
config firewall policy
   edit 0
      set srcintf port1
      set dstintf port2
      set srcaddr all
      set dstaddr all
      set action accept
      set service ALL
      set schedule always
      set wanopt enable
      set wanopt-profile default
      set wanopt-detection off
      set wanopt-peer Server-Fgt
   end
```

To configure the server-side FortiGate unit

1. Add the Local Host ID to the server-side FortiGate configuration:

```
config wanopt settings
   set host-id Server-Fgt
end
```

2. Add the client-side Local Host ID to the server-side peer list:

```
config wanopt peer
   edit Client-Fgt
   set ip 192.168.30.12
end
```

3. Add a WAN optimization tunnel explicit proxy policy.

```
configure firewall explicit-proxy-policy
   edit 0
      set proxy wanopt
      set dstintf port1
      set srcaddr all
      set dstaddr all
```

```
        set action accept
        set schedule always
        set service ALL
    next
end
```

Testing and troubleshooting the configuration

To test the configuration attempt to start a web browsing session between the client network and the web server network. For example, from a PC on the client network browse to the IP address of a web server on the web server network, for example http://192.168.10.100. Even though this address is not on the client network you should be able to connect to this web server over the WAN optimization tunnel.

If you can connect, check WAN optimization monitoring. If WAN optimization has been forwarding the traffic the WAN optimization monitor should show the protocol that has been optimized (in this case HTTP) and the reduction rate in WAN bandwidth usage.

If you can't connect you can try the following to diagnose the problem:

- Review your configuration and make sure all details such as address ranges, peer names, and IP addresses are correct.
- Confirm that the security policy on the client-side FortiGate unit is accepting traffic for the 192.168.10.0 network. You can do this by checking the policy monitor (**Monitor > Firewall Monitor**). Look for sessions that use the policy ID of this policy.
- Check routing on the FortiGate units and on the client and web server networks to make sure packets can be forwarded as required. The FortiGate units must be able to communicate with each other, routing on the client network must allow packets destined for the web server network to be received by the client-side FortiGate unit, and packets from the server-side FortiGate unit must be able to reach the web servers.

You can use the following `get` and `diagnose` commands to display information about how WAN optimization is operating.

Enter the following command to list all of the running WAN optimization tunnels and display information about each one. The command output for the client-side FortiGate unit shows 10 tunnels all created by peer-to-peer WAN optimization rules (auto-detect set to off).

```
diagnose wad tunnel list

Tunnel: id=100 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web-servers id=100 ip=192.168.30.12
   SSL-secured-tunnel=no auth-grp=
   bytes_in=348 bytes_out=384

Tunnel: id=99 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web-servers id=99 ip=192.168.30.12
SSL-secured-tunnel=no auth-grp=
bytes_in=348 bytes_out=384

Tunnel: id=98 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web-servers id=98 ip=192.168.30.12
   SSL-secured-tunnel=no auth-grp=
   bytes_in=348 bytes_out=384
```

```
Tunnel: id=39 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web-servers id=39 ip=192.168.30.12
   SSL-secured-tunnel=no auth-grp=
   bytes_in=1068 bytes_out=1104

Tunnel: id=7 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web-servers id=7 ip=192.168.30.12
   SSL-secured-tunnel=no auth-grp=
   bytes_in=1228 bytes_out=1264

Tunnel: id=8 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web-servers id=8 ip=192.168.30.12
   SSL-secured-tunnel=no auth-grp=
   bytes_in=1228 bytes_out=1264

Tunnel: id=5 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web-servers id=5 ip=192.168.30.12
   SSL-secured-tunnel=no auth-grp=
   bytes_in=1228 bytes_out=1264

Tunnel: id=4 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web-servers id=4 ip=192.168.30.12
   SSL-secured-tunnel=no auth-grp=
   bytes_in=1228 bytes_out=1264

Tunnel: id=1 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web-servers id=1 ip=192.168.30.12
   SSL-secured-tunnel=no auth-grp=
   bytes_in=1228 bytes_out=1264

Tunnel: id=2 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web-servers id=2 ip=192.168.30.12
   SSL-secured-tunnel=no auth-grp=
   bytes_in=1228 bytes_out=1264

Tunnels total=10 manual=10 auto=0
```

Example Active-passive WAN optimization

In active-passive WAN optimization you add an active WAN optimization policy to the client-side FortiGate unit and you add a WAN optimization tunnel policy and a passive WAN optimization policy to the server-side FortiGate unit.

The active policy accepts the traffic to be optimized and sends it down the WAN optimization tunnel to the server-side FortiGate unit. The active policy can also apply security profiles and other features to traffic before it exits the client-side FortiGate unit.

A tunnel explicit proxy policy on the sever-side FortiGate unit allows the server-side FortiGate unit to form a WAN optimization tunnel with the client-side FortiGate unit. The passive WAN optimization policy is required because of the active policy on the client-side FortiGate unit. You can also use the passive policy to apply WAN optimization transparent mode and features such as security profiles, logging, traffic shaping and web caching to the traffic before it exits the server-side FortiGate unit.

Network topology and assumptions

On the client-side FortiGate unit this example configuration includes a WAN optimization profile that optimizes CIFS, HTTP, and FTP traffic and an active WAN optimization policy. The active policy also applies virus scanning to the WAN optimization traffic.

On the server-side FortiGate unit, the passive policy applies application control to the WAN optimization traffic.

In this example, WAN optimization transparent mode is selected in the WAN optimization profile and the passive WAN optimization policy accepts this transparent mode setting. This means that the optimized packets maintain their original source and destination addresses. As a result, routing on the client network must be configured to route packets for the server network to the client-side FortiGate unit. Also the routing configuration on the server network must be able to route packets for the client network to the server-side FortiGate unit.

Example active-passive WAN optimization topology

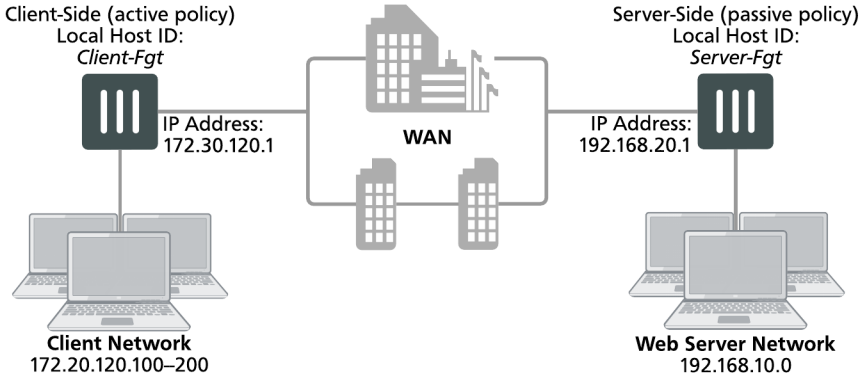

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Configure the client-side FortiGate unit:
   - Add peers.
   - Add a WAN optimization profile to optimize CIFS, FTP, and HTTP traffic.
   - Add firewall addresses for the client and web server networks.
   - Add an active WAN optimization policy.
2. Configure the server-side FortiGate unit by:
   - Add peers.
   - Add firewall addresses for the client and web server networks.

- Add a passive WAN optimization policy.
- Add a WAN optimization tunnel policy.

Configuring basic active-passive WAN optimization - web-based manager

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit web-based manager.

To configure the client-side FortiGate unit

1. Go to **WAN Opt. & Cache > Peers** and enter a **Local Host ID** for the client-side FortiGate unit:

Local Host ID	Client-Fgt

2. Select **Apply**.
3. Select **Create New** and add a Peer Host ID and the **IP Address** for the server-side FortiGate unit:

Peer Host ID	Server-Fgt
IP Address	192.168.20.1

4. Select **OK**.
5. Go to **WAN Opt. & Cache > Profiles** and select Create New to add a WAN optimization profile to optimize CIFS, HTTP, and FTP traffic:

Name	Custom-wan-opt-pro
Transparent Mode	Select

6. Select the **CIFS** protocol, select **Byte Caching** and set the **Port** to 445.
7. Select the **FTP** protocol, select **Byte Caching** and set the **Port** to 21.
8. Select the **HTTP** protocol, select **Byte Caching** and set the **Port** to 80.
9. Select **OK**.
10. Go to **Policy & Objects > Addresses** and select **Create New** to add an address for the client network.

Category	Address
Address Name	Client-Net
Type	IP Range
Subnet / IP Range	172.20.120.100-172.20.120.200
Interface	port1

11. Select **Create New** to add an address for the web server network.

Category	Address

Address Name	Web-Server-Net
Type	Subnet
Subnet / IP Range	192.168.10.0/24
Interface	port2

12. Go to **Policy & Objects > IPv4 Policy** and select **Create New** to add an active WAN optimization security policy:

Incoming Interface	port1
Source Address	Client-Net
Outgoing Interface	port2
Destination Address	Web-Server-Net
Schedule	always
Service	HTTP FTP SMB
Action	ACCEPT

13. Turn on **WAN Optimization** and configure the following settings:

WAN Optimization	active
Profile	Custom-wan-opt-pro

14. Turn on Antivirus and select the **default** antivirus profile.
15. Select **OK**.

To configure the server-side FortiGate unit

1. Go to **WAN Opt. & Cache > Peers** and enter a **Local Host ID** for the server-side FortiGate unit:

Local Host ID	Server-Fgt

2. Select **Apply**.
3. Select **Create New** and add a **Peer Host ID** and the **IP Address** for the client-side FortiGate unit:

Peer Host ID	Client-Fgt
IP Address	172.30.120.1

4. Select **OK**.
5. Go to **Policy & Objects > Addresses** and select **Create New** to add an address for the client network.

Category	Address
Address Name	Client-Net
Type	IP Range
Subnet / IP Range	172.20.120.100-172.20.120.200
Interface	port1

6. Select **Create New** to add a firewall address for the web server network.

Category	Address
Address Name	Web-Server-Net
Type	Subnet
Subnet / IP Range	192.168.10.0/24
Interface	port2

7. Select **OK**.
8. Select **Policy & Objects > IPv4 Policy** and select **Create New** to add a passive WAN optimization policy that applies application control.

Incoming Interface	port2
Source Address	Client-Net
Outgoing Interface	port1
Destination Address	Web-Server-Net
Schedule	always
Service	ALL
Action	ACCEPT

9. Turn on **WAN Optimization** and configure the following settings:

WAN Optimization	passive
Passive Option	default

10. Select **OK**.
11. From the CLI enter the following command to add a WAN optimization tunnel explicit proxy policy.

```
configure firewall explicit-proxy-policy
   edit 0
      set proxy wanopt
      set dstintf port1
      set srcaddr all
      set dstaddr all
```

```
      set action accept
      set schedule always
      set service ALL
   next
end
```

Configuring basic active-passive WAN optimization - CLI

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit CLI.

To configure the client-side FortiGate unit

1. Add the Local Host ID to the client-side FortiGate configuration:

```
config wanopt settings
   set host-id Client-Fgt
end
```

2. Add the server-side Local Host ID to the client-side peer list:

```
config wanopt peer
   edit Server-Fgt
   set ip 192.168.20.1
end
```

3. Add a WAN optimization profile to optimize CIFS, HTTP, and FTP traffic.

```
config wanopt profile
   edit Custom-wan-opt-pro
      config cifs
         set status enable
         set byte-caching enable
         set port 445
      end
      config http
         set status enable
         set byte-caching enable
         set port 80
      end
      config ftp
         set status enable
         set byte-caching enable
         set port 21
      end
   end
```

4. Add a firewall address for the client network.

```
config firewall address
   edit Client-Net
      set type iprange
      set start-ip 172.20.120.100
      set end-ip 172.20.120.200
      set associated-interface port1
   end
```

5. Add a firewall address for the web server network.

```
config firewall address
   edit Web-Server-Net
      set type ipmask
      set subnet 192.168.10.0 255.255.255.0
      set associated-interface port2
```

```
   end
```

6. Add an active WAN optimization security policy that applies virus scanning:

```
config firewall policy
   edit 0
      set srcintf port1
      set dstintf port2
      set srcaddr Client-net
      set dstaddr Web-Server-Net
      set action accept
      set service HTTP FTP SMB
      set schedule always
      set wanopt enable
      set wanopt-detection active
      set wanopt-profile Custom-wan-opt-pro
   end
```

To configure the server-side FortiGate unit

1. Add the Local Host ID to the server-side FortiGate configuration:

```
config wanopt settings
   set host-id Server-Fgt
end
```

2. Add the client-side Local Host ID to the server-side peer list:

```
config wanopt peer
   edit Client-Fgt
   set ip 172.20.120.1
end
```

3. Add a firewall address for the client network.

```
config firewall address
   edit Client-Net
      set type iprange
      set start-ip 172.20.120.100
      set end-ip 172.20.120.200
      set associated-interface port1
   end
```

4. Add a firewall address for the web server network.

```
config firewall address
   edit Web-Server-Net
      set type ipmask
      set subnet 192.168.10.0 255.255.255.0
      set associated-interface port2
   end
```

5. Add a passive WAN optimization policy.

```
config firewall policy
   edit 0
      set srcintf port1
      set dstintf port2
      set srcaddr Client-Net
      set dstaddr Web-Server-Net
      set action accept
      set service ALL
      set schedule always
      set wanopt enable
      set wanopt-detection passive
      set wanopt-passive-opt default
```

```
    end
```

6. Add a WAN optimization tunnel explicit proxy policy.

```
configure firewall explicit-proxy-policy
   edit 0
      set proxy wanopt
      set dstintf port1
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service ALL
   next
end
```

Testing and troubleshooting the configuration

To test the configuration attempt to start a web browsing session between the client network and the web server network. For example, from a PC on the client network browse to the IP address of a web server on the web server network, for example http://192.168.10.100. Even though this address is not on the client network you should be able to connect to this web server over the WAN optimization tunnel.

If you can connect, check WAN optimization monitoring. If WAN optimization has been forwarding the traffic the WAN optimization monitor should show the protocol that has been optimized (in this case HTTP) and the reduction rate in WAN bandwidth usage.

If you can't connect you can try the following to diagnose the problem:

- Review your configuration and make sure all details such as address ranges, peer names, and IP addresses are correct.
- Confirm that the security policy on the Client-Side FortiGate unit is accepting traffic for the 192.168.10.0 network and that this security policy does not include security profiles. You can do this by checking the FortiGate session table from the dashboard. Look for sessions that use the policy ID of this policy.
- Check routing on the FortiGate units and on the client and web server networks to make sure packets can be forwarded as required. The FortiGate units must be able to communicate with each other, routing on the client network must allow packets destined for the web server network to be received by the client-side FortiGate unit, and packets from the server-side FortiGate unit must be able to reach the web servers etc.

You can use the following `get` and `diagnose` commands to display information about how WAN optimization is operating

Enter the following command to list all of the running WAN optimization tunnels and display information about each one. The command output shows 3 tunnels all created by peer-to-peer WAN optimization rules (auto-detect set to on).

```
diagnose wad tunnel list

Tunnel: id=139 type=auto
   vd=0 shared=no uses=0 state=1
   peer name= id=0 ip=unknown
   SSL-secured-tunnel=no auth-grp=test
   bytes_in=744 bytes_out=76

Tunnel: id=141 type=auto
   vd=0 shared=no uses=0 state=1
   peer name= id=0 ip=unknown
```

```
   SSL-secured-tunnel=no auth-grp=test
   bytes_in=727 bytes_out=76

Tunnel: id=142 type=auto
   vd=0 shared=no uses=0 state=1
   peer name= id=0 ip=unknown
   SSL-secured-tunnel=no auth-grp=test
   bytes_in=727 bytes_out=76

Tunnels total=3 manual=0 auto=3
```

Example Adding secure tunneling to an active-passive WAN optimization configuration

This example shows how to configure two FortiGate units for active-passive WAN optimization with secure tunneling. The same authentication group is added to both FortiGate units. The authentication group includes a password (or pre-shared key) and has **Peer Acceptance** set to **Accept any Peer**. An active policy is added to the client-side FortiGate unit and a passive policy to the server-side FortiGate unit. The active policy includes a profile that performs secure tunneling, optimizes HTTP traffic, and uses Transparent Mode and byte caching.

The authentication group is named **Auth-Secure-Tunnel** and the password for the pre-shared key is **2345678**. The topology for this example is shown below. This example includes web-based manager configuration steps followed by equivalent CLI configuration steps. For information about secure tunneling, see Secure tunneling on page 49.

Network topology and assumptions

This example configuration includes a client-side FortiGate unit called Client-net with a WAN IP address of 172.30.120.1.This unit is in front of a network with IP address 172.20.120.0. The server-side FortiGate unit is called Web-servers and has a WAN IP address of 192.168.20.1. This unit is in front of a web server network with IP address 192.168.10.0.

Example active-passive WAN optimization and secure tunneling topology

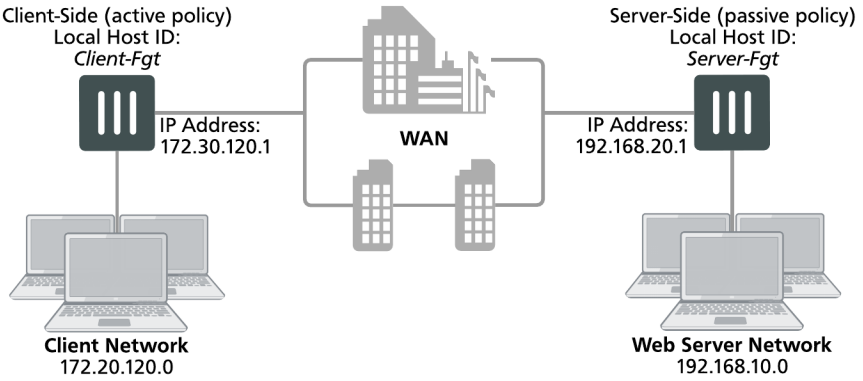

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Configure the client-side FortiGate unit:
   - Add peers.
   - Add an authentication group.
   - Add an active WAN optimization policy.
2. Configure the server-side FortiGate unit.
   - Add peers.
   - Add the same authentication group
   - Add a passive WAN optimization policy that applies application control.
   - Add a WAN optimization tunnel policy.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

Configuring WAN optimization with secure tunneling - web-based manager

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit web-based manager. (CLI steps follow.)

To configure the client-side FortiGate unit

1. Go to **WAN Opt. & Cache > Peers** and enter a **Local Host ID** for the client-side FortiGate unit:

Local Host ID	Client-Fgt

2. Select **Apply** to save your setting.
3. Select **Create New** and add a **Peer Host ID** and the **IP Address** for the server-side FortiGate unit:

Peer Host ID	Server-Fgt
IP Address	192.168.20.1

4. Select **OK**.
5. Go to **WAN Opt. & Cache > Authentication Groups** and select **Create New** to add the authentication group to be used for secure tunneling:

Name	Auth-Secure-Tunnel
Authentication Method	Pre-shared key
Password	2345678
Peer Acceptance	Accept Any Peer

6. Select **OK**.
7. Go to **WAN Opt. & Cache > Profiles** and select **Create New** to add a WAN optimization profile that enables secure tunneling and includes the authentication group:

Name	Secure-wan-op-pro
Transparent Mode	Select
Authentication Group	Auth-Secure-tunnel

8. Select the **HTTP** protocol, select Secure Tunneling and **Byte Caching** and set the **Port** to 80.
9. Select **OK**.
10. Go to **Policy & Objects > Addresses** and select **Create New** to add a firewall address for the client network.

Category	Address
Name	Client-Net
Type	Subnet
Subnet / IP Range	172.20.120.0/24
Interface	port1

11. Select **Create New** to add a firewall address for the web server network.

Category	Address
Address Name	Web-Server-Net
Type	Subnet
Subnet / IP Range	192.168.10.0/24
Interface	port2

12. Go to **Policy & Objects > IPv4 Policy** and select **Create New** to add an active WAN optimization security policy:

Incoming Interface	port1
Source Address	Client-Net
Outgoing Interface	port2
Destination Address	Web-Server-Net
Schedule	always
Service	HTTP
Action	ACCEPT

13. Turn on **WAN Optimization** and configure the following settings:

WAN Optimization	active
Profile	Secure-wan-opt-pro

14. Select **OK**.

To configure the server-side FortiGate unit

1. Go to **WAN Opt. & Cache > Peers** and enter a **Local Host ID** for the server-side FortiGate unit:

Local Host ID	Server-Fgt

2. Select **Apply** to save your setting.
3. Select **Create New** and add a **Peer Host ID** and the **IP Address** for the client-side FortiGate unit:

Peer Host ID	Client-Fgt
IP Address	172.30.120.1

4. Select **OK**.
5. Go to **WAN Opt. & Cache > Authentication Groups** and select **Create New** and add an authentication group to be used for secure tunneling:

Name	Auth-Secure-Tunnel
Authentication Method	Pre-shared key
Password	2345678
Peer Acceptance	Accept Any Peer

6. Select **OK**.
7. Go to **Policy & Objects > Addresses** and select **Create New** to add a firewall address for the client network.

Category	Address
Name	Client-Net
Type	Subnet
Subnet / IP Range	172.20.120.0/24
Interface	port1

8. Select **Create New** to add a firewall address for the web server network.

Category	Address

Address Name	Web-Server-Net
Type	Subnet
Subnet / IP Range	192.168.10.0/24
Interface	port2

9. Select **OK**.
10. Select **Create New** to add a passive WAN optimization policy that applies application control.

Incoming Interface	port2
Source Address	Client-Net
Outgoing Interface	port1
Destination Address	Web-Server-Net
Schedule	always
Service	ALL
Action	ACCEPT

11. Turn on **WAN Optimization** and configure the following settings:

WAN Optimization	passive
Passive Option	default

12. Select **OK**.
13. From the CLI enter the following command to add a WAN optimization tunnel explicit proxy policy.

```
configure firewall explicit-proxy-policy
   edit 0
      set proxy wanopt
      set dstintf port1
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service ALL
   next
end
```

Configuring WAN optimization with secure tunneling - CLI

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit CLI.

To the client-side FortiGate unit

1. Add the Local Host ID to the client-side FortiGate configuration:

```
config wanopt settings
   set host-id Client-Fgt
end
```

2. Add the server-side Local Host ID to the client-side peer list:

```
config wanopt peer
   edit Server-Fgt
   set ip 192.168.20.1
end
```

3. Add a new authentication group to be used for secure tunneling:

```
config wanopt auth-group
   edit Auth-Secure-Tunnel
      set auth-method psk
      set psk 2345678
end
```

   Leave `peer-accept` at its default value.

4. Add a WAN optimization profile that enables secure tunneling and includes the authentication group, enables HTTP protocol optimization, and enables secure tunneling and byte caching for HTTP traffic:

```
config wanopt profile
   edit Secure-wan-op-pro
      set auth-group Auth-Secure-Tunnel
      config http
         set status enable
         set secure-tunnel enable
         set byte-caching enable
         set port 80
      end
   end
```

5. Add a firewall address for the client network.

```
config firewall address
   edit Client-Net
      set type ipmask
      set subnet 172.20.120.0 255.255.255.0
      set associated-interface port1
   end
```

6. Add a firewall address for the web server network.

```
config firewall address
   edit Web-Server-Net
      set type ipmask
      set subnet 192.168.10.0 255.255.255.0
      set associated-interface port2
   end
```

7. Add an active WAN optimization security policy that includes the WAN optimization profile that enables secure tunneling and that applies virus scanning:

```
config firewall policy
   edit 0
      set srcintf port1
      set dstintf port2
      set srcaddr Client-Net
      set dstaddr Web-Server-Net
      set action accept
```

```
      set service HTTP
      set schedule always
      set wanopt enable
      set wanopt-detection active
      set wanopt-profile Secure-wan-opt-pro
   end
```

To configure the server-side FortiGate unit

1. Add the Local Host ID to the server-side FortiGate configuration:

```
config wanopt settings
   set host-id Server-Fgt
end
```

2. Add the client-side Local Host ID to the server-side peer list:

```
config wanopt peer
   edit Client-Fgt
   set ip 172.20.120.1
end
```

3. Add an authentication group to be used for secure tunneling:

```
config wanopt auth-group
   edit Auth-Secure-Tunnel
      set auth-method psk
      set psk 2345678
   end
```

   Leave `peer-accept` at its default value.

4. Add a firewall address for the client network.

```
config firewall address
   edit Client-Net
      set type ipmask
      set subnet 172.20.120.0 255.255.255.0
      set associated-interface port1
   end
```

5. Add a firewall address for the web server network.

```
config firewall address
   edit Web-Server-Net
      set type ipmask
      set subnet 192.168.10.0 255.255.255.0
      set associated-interface port2
   end
```

6. Add a passive WAN optimization policy.

```
config firewall policy
   edit 0
      set srcintf port1
      set dstintf port2
      set srcaddr Client-Net
      set dstaddr Web-Server-Net
      set action accept
      set service ALL
      set schedule always
      set wanopt enable
```

```
      set wanopt-detection passive
      set wanopt-passive-opt default
   end
```

7. Add a WAN optimization tunnel explicit proxy policy.

```
configure firewall explicit-proxy-policy
   edit 0
      set proxy wanopt
      set dstintf port1
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service ALL
   next
end
```

Web caching and SSL offloading

FortiGate web caching is a form of object caching that accelerates web applications and web servers by reducing bandwidth usage, server load, and perceived latency. Web caching supports caching of HTTP 1.0 and HTTP 1.1 web sites. See RFC 2616 for information about web caching for HTTP 1.1.

Web caching supports caching of Flash content over HTTP but does not cache audio and video streams including Flash videos and streaming content that use native streaming protocols such as RTMP.

The first time a file is received by web caching it is cached in the format it is received in, whether it be compressed or uncompressed. When the same file is requested by a client but in a different compression format, the cached file is converted to the new compressed format before being sent to the client.

There are three significant advantages to using web caching to improve HTTP and WAN performance:

- reduced bandwidth consumption because fewer requests and responses go over the WAN or Internet.
- reduced web server load because there are fewer requests for web servers to handle.
- reduced latency because responses for cached requests are available from a local FortiGate unit instead of from across the WAN or Internet.

You can use web caching to cache any web traffic that passes through the FortiGate unit, including web pages from web servers on a LAN, WAN or on the Internet. You apply web caching by enabling the web caching option in any security policy. When enabled in a security policy, web caching is applied to all HTTP sessions accepted by the security policy. If the security policy is an explicit web proxy security policy, the FortiGate unit caches explicit web proxy sessions.

Turning on web caching for HTTP and HTTPS traffic

Web caching can be applied to any HTTP or HTTPS traffic by enabling web caching in a security policy that accepts the traffic. This includes IPv4, IPv6, WAN optimization and explicit web proxy traffic. Web caching caches all HTTP traffic accepted by a policy on TCP port 80.

You can add web caching to a policy to:

- Cache Internet HTTP traffic for users on an internal network to reduce Internet bandwidth use. Do this by selecting the web cache option for security policies that allow users on the internal network to browse web sites on the Internet.
- Reduce the load on a public facing web server by caching objects on the FortiGate unit. This is a reverse proxy with web caching configuration. Do this by selecting the web cache option for a security policy that allows users on the Internet to connect to the web server.
- Cache outgoing explicit web proxy traffic when the explicit proxy is used to proxy users in an internal network who are connecting to the web servers on the Internet. Do this by selecting the web cache option for explicit web proxy security policies that allow users on the internal network to browse web sites on the Internet.
- Combine web caching with WAN optimization. You can enable web caching in any WAN optimization security policy. This includes manual, active, and passive WAN optimization policies and WAN optimization tunnel policies.

You can enable web caching on both the client-side and the server-side FortiGate units or on just one or the other. For optimum performance you can enable web caching on both the client-side and server-side FortiGate units. In this way only uncached content is transmitted through the WAN optimization tunnel. All cached content is access locally by clients from the client side FortiGate unit.

One important use for web caching is to cache software updates (for example, Windows Updates or iOS updates. When updates occur a large number of users may all be trying to download these updates at the same time. Caching these updates will be a major performance improvement and also have a potentially large impact on reducing Internet bandwidth use. You may want to adjust the maximum cache object size to make sure these updates are cached. See Max cache object size on page 78.

Turning on web caching for HTTPS traffic

Web caching can also cache the content of HTTPS traffic on TCP port 443. With HTTPS web caching, the FortiGate unit receives the HTTPS traffic on behalf of the client, opens up the encrypted traffic and extracts content to be cached. Then FortiGate unit re-encrypts the traffic and sends it on to its intended recipient. It is very similar to a man-in-the-middle attack.

You enable HTTPS web caching from the CLI in a security policy or an explicit proxy policy that accepts the traffic to be cached using `webcache-https`. For a firewall policy:

```
config firewall policy
   edit 0
      .
      .
      .
      set webcache enable
      set webcache-https any
      .
      .
      .
   end
```

For an explicit web proxy policy:

```
config firewall policy
   edit 0
      set proxy web
      .
      .
      .
      set webcache enable
      set webcache-https any
      .
      .
      .
   end
```

Web caching for HTTPS traffic is not supported if WAN optimization is enabled.

The `any` setting causes the FortiGate unit to re-encrypt the traffic with the FortiGate unit's certificate rather than the original certificate. This configuration can cause errors for HTTPS clients because the name on the certificate does not match the name on the web site.

You can stop these errors from happening by configuring HTTPS web caching to use the web server's certificate by setting `webcache-https` to `ssl-server`. This option is available for both firewall policies and explicit web proxy policies.

```
config firewall policy
   edit 0
      .
      .
      .
      set webcache enable
      set webcache-https ssl-server
      .
      .
      .
   end
```

The `ssl-server` option causes the FortiGate unit to re-encrypt the traffic with a certificate that you imported into the FortiGate unit. You can add certificates using the following command:

```
config firewall ssl-server
   edit corporate-server
      set ip <Web-Server-IP>
      set port 443
      set ssl-mode { full | half}
      set ssl-cert <Web-Server-Cert>
   end
```

Where:

`Web-Server-IP` is the web server's IP address.

`Web-Server-Cert` is a web server certificate imported into the FortiGate unit.

The SSL server configuration also determines whether the SSL server is operating in half or full mode and the port used for the HTTPS traffic.

You can add multiple SSL server certificates in this way. When web caching processing an SSL stream if it can find a certificate that matches the web server IP address and port of one of the added SSL servers; that certificate is used to encrypt the SSL traffic before sending it to the client. As a result the client does not generate SSL certificate errors.

Web caching uses the FortiGate unit's FortiASIC to accelerate SSL decryption/encryption performance.

Full mode SSL server configuration

The `ssl-mode` option determines whether the SSL server operates in half or full mode. In full mode the FortiGate unit performs both decryption and encryption of the HTTPS traffic. The full mode sequence is shown below.

Full mode SSL server configuration

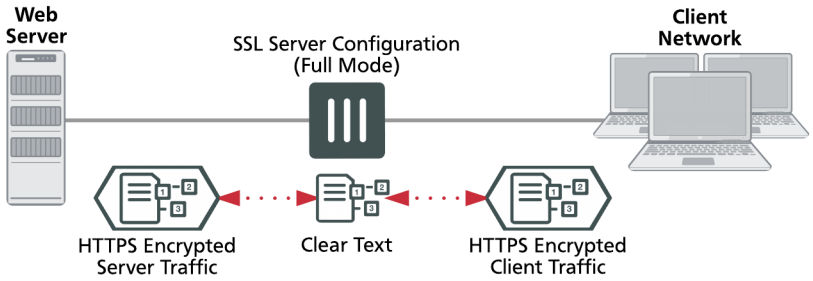

In full mode the FortiGate unit is acting as a man in the middle, decrypting and encrypting the traffic. So both the client and the web server see encrypted packets.

Usually the port of the encrypted HTTPS traffic is always 443. However, in the SSL server configuration you can set the port used for HTTPS traffic. This port is not altered by the SSL Server. So for example, if the SSL Server receives HTTPS traffic on port 443, the re-encrypted traffic forwarded to the FortiGate unit to the server or client will still use port 443.

Half mode SSL server configuration

In half mode, the FortiGate unit only performs one encryption or decryption action. If HTTP packets are received, the half mode SSL server encrypts them and converts them to HTTPS packets. If HTTPS packets are received, the SSL server decrypts them and converts them to HTTP packets.

Half mode SSL server configuration

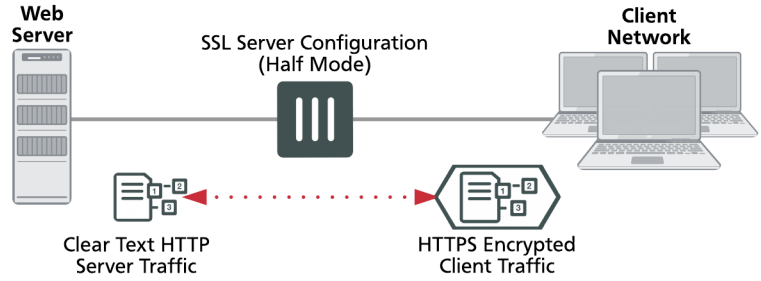

In half mode, the FortiGate unit is acting like an SSL accelerator, offloading HTTPS decryption from the web server to the FortiGate unit. Since FortiGate units can accelerate SSL processing, the end result could be improved web site performance.

Usually the port of the encrypted traffic is always 443. However, in the SSL server configuration you can set the port used for HTTPS traffic. No matter what port is used for the HTTPS traffic, the decrypted HTTP traffic uses port 80.

Changing the ports on which to look for HTTP and HTTPS traffic to cache

By default FortiOS assumes HTTP traffic uses TCP port 80 and HTTPS traffic uses port 443. So web caching caches all HTTP traffic accepted by a policy on TCP port 80 and all HTTPS traffic on TCP port 443. If you want to cache HTTP or HTTPS traffic on other ports, you can enable security profiles for the security policy and configure a proxy options profile to that looks for HTTP and HTTPS traffic on other TCP ports. To configure a proxy options profile go to **Network > Explicit Proxy**.

Setting the HTTP port to **Any** in a proxy options profile is not compatible with web caching. If you set the HTTP port to any, web caching only caches HTTP traffic on port 80.

Web caching and HA

You can configure web caching on a FortiGate HA cluster. The recommended best practice HA configuration for web caching is active-passive mode. When the cluster is operating, all web caching sessions are processed by the primary unit only. Even if the cluster is operating in active-active mode, HA does not load-balance web caching sessions.

In a cluster, only the primary unit stores the web cache database. The databases is not synchronized to the subordinate units. So, after a failover, the new primary unit must build its web cache.

Web caching and memory usage

To accelerate and optimize disk access and to provide better throughput and less latency, web caching uses provisioned memory to reduce disk I/O and increase disk I/O efficiency. In addition, web caching requires a small amount of additional memory per session for comprehensive flow control logic and efficient traffic forwarding.

When web caching is enabled you will see a reduction in available memory. The reduction increases when more web caching sessions are being processed. If you are thinking of enabling web caching on an operating FortiGate unit, make sure its memory usage is not maxed out during high traffic periods.

In addition to using the system dashboard to see the current memory usage you can use the `get test wad 2` command to see how much memory is currently being used by web caching. See get test {wad | wccpd} <test_ level> on page 145 for more information.

Changing web cache settings

In most cases, the default settings for the WAN optimization web cache are acceptable. However, you may want to change them to improve performance or optimize the cache for your configuration. To change these settings, go to **WAN Opt. & Cache > Settings.**

From the FortiGate CLI, you can use the `config wanopt webcache` command to change these WAN optimization web cache settings.

For more information about many of these web cache settings, see RFC 2616.

Always revalidate

Select to always revalidate requested cached objects with content on the server before serving them to the client.

Max cache object size

Set the maximum size of objects (files) that are cached. The default size is 512000 KB and the range is 1 to 4294967 KB. This setting determines the maximum object size to store in the web cache. Objects that are larger than this size are still delivered to the client but are not stored in the FortiGate web cache.

For most web traffic the default maximum cache object size is recommended. However, since web caching can also cache larger objects such as Windows updates, Mac OS updates, iOS updates or other updates delivered using HTTP you might want to increase the object size to make sure these updates are cached. Caching these updates can save a lot of Internet bandwidth and improve performance when major updates are released by these vendors.

Negative response duration

Set how long in minutes that the FortiGate unit caches error responses from web servers. If error responses are cached, then subsequent requests to the web cache from users will receive the error responses regardless of the actual object status.

The default is 0, meaning error responses are not cached. The content server might send a client error code (4xx HTTP response) or a server error code (5xx HTTP response) as a response to some requests. If the web cache is configured to cache these negative responses, it returns that response in subsequent requests for that page or image for the specified number of minutes.

Fresh factor

Set the fresh factor as a percentage. The default is 100, and the range is 1 to 100%. For cached objects that do not have an expiry time, the web cache periodically checks the server to see if the objects have expired. The higher the **Fresh Factor** the less often the checks occur.

For example, if you set the **Max TTL** value and **Default TTL** to 7200 minutes (5 days) and set the **Fresh Factor** to 20, the web cache check the cached objects 5 times before they expire, but if you set the **Fresh Factor** to 100, the web cache will check once.

Max TTL

The maximum amount of time (Time to Live) an object can stay in the web cache without the cache checking to see if it has expired on the server. The default is 7200 minutes (120 hours or 5 days) and the range is 1 to 5256000 minutes (5256000 minutes in a year).

Min TTL

The minimum amount of time an object can stay in the web cache before the web cache checks to see if it has expired on the server. The default is 5 minutes and the range is 1 to 5256000 minutes (5256000 minutes in a year).

Default TTL

The default expiry time for objects that do not have an expiry time set by the web server. The default expiry time is 1440 minutes (24 hours) and the range is 1 to 5256000 minutes (5256000 minutes in a year).

Proxy FQDN

The fully qualified domain name (FQDN) for the proxy server. This is the domain name to enter into browsers to access the proxy server. This field is for information only can be changed from the explicit web proxy configuration.

Max HTTP request length

The maximum length of an HTTP request that can be cached. Larger requests will be rejected. This field is for information only can be changed from the explicit web proxy configuration.

Max HTTP message length

The maximum length of an HTTP message that can be cached. Larger messages will be rejected. This field is for information only can be changed from the explicit web proxy configuration.

Ignore

Select the following options to ignore some web caching features.

If-modified-since	By default, if the time specified by the if-modified-since (IMS) header in the client's conditional request is greater than the last modified time of the object in the cache, it is a strong indication that the copy in the cache is stale. If so, HTTP does a conditional GET to the Overlay Caching Scheme (OCS), based on the last modified time of the cached object. Enable ignoring if-modified-since to override this behavior.
HTTP 1.1 conditionals	HTTP 1.1 provides additional controls to the client over the behavior of caches toward stale objects. Depending on various cache-control headers, the FortiGate unit can be forced to consult the OCS before serving the object from the cache. For more information about the behavior of cache-control header values, see RFC 2616.Enable ignoring HTTP 1.1 Conditionals to override this behavior.

Pragma-no-cache	Typically, if a client sends an HTTP GET request with a pragma no-cache (PNC) or cache-control no-cache header, a cache must consult the OCS before serving the content. This means that the FortiGate unit always re-fetches the entire object from the OCS, even if the cached copy of the object is fresh.Because of this behavior, PNC requests can degrade performance and increase server-side bandwidth utilization. However, if you enable ignoring Pragma-no-cache, then the PNC header from the client request is ignored. The FortiGate unit treats the request as if the PNC header is not present.
IE Reload	Some versions of Internet Explorer issue Accept / header instead of Pragma no-cache header when you select **Refresh**. When an Accept header has only the / value, the FortiGate unit treats it as a PNC header if it is a type-N object. Enable ignoring IE reload to cause the FortiGate unit to ignore the PNC interpretation of the Accept / header.

Cache Expired Objects

Applies only to type-1 objects. When this option is selected, expired type-1 objects are cached (if all other conditions make the object cacheable).

Revalidated Pragma-no-cache

The pragma-no-cache (PNC) header in a client's request can affect how efficiently the FortiGate unit uses bandwidth. If you do not want to completely ignore PNC in client requests (which you can do by selecting to ignore Pragma-no-cache, above), you can nonetheless lower the impact on bandwidth usage by selecting **Revalidate Pragma-no-cache**.

When you select **Revalidate Pragma-no-cache**, a client's non-conditional PNC-GET request results in a conditional GET request sent to the OCS if the object is already in the cache. This gives the OCS a chance to return the 304 Not Modified response, which consumes less server-side bandwidth, because the OCS has not been forced to otherwise return full content.

By default, **Revalidate Pragma-no-cache** is disabled and is not affected by changes in the top-level profile.

Most download managers make byte-range requests with a PNC header. To serve such requests from the cache, you should also configure byte-range support when you configure the **Revalidate pragma-no-cache** option.

Forwarding URLs to forwarding servers and exempting web sites from web caching

You can go to **Network > Explicit Proxy** and use the URL match list to forward URL patterns to forwarding servers and create a list of URLs that are exempt from web caching.

Forwarding URLs and URL patterns to forwarding servers

As part of configuring the explicit web proxy you can configure proxy chaining by adding web proxy forwarding servers. See "Proxy chaining (web proxy forwarding servers) ".

You can then use the URL match list to always forward explicit web proxy traffic destined for configured URLs or URL patterns to one of these forwarding servers. For example, you might want to forward all traffic for a specific country to a proxy server located in that country.

To forward traffic destined for a URL to a forwarding server that you have already added, go to **Network > Explicit Proxy** and select **Create New**. Add a name for the URL match entry and enter the URL or URL pattern. You can use wildcards such as * and ? and you can use a numeric IP address. Select **Forward to Server** and select a web proxy forwarding server from the list.

You can also exempt the URL or URL pattern from web caching.

Use the following command to forward all .ca traffic to a proxy server and all .com traffic to another proxy server.

```
config web-proxy url-match
   edit "com"
      set forward-server "server-commercial"
      set url-pattern "com"
   next
   edit "ca"
      set forward-server "server-canada"
      set url-pattern "ca"
   next
   edit "www.google.ca"
      set cache-exemption enable
      set url-pattern "www.google.ca"
   next
end
```

Exempting web sites from web caching

You may want to exempt some URLs from web caching for a number of reasons. For example, if your users access websites that are not compatible with FortiGate web caching you can add the URLs of these web sites to the web caching exempt list. You can add URLs and numeric IP addresses to the web cache exempt list.

You can also add URLs to the web cache exempt list by going to **Network > Explicit Proxy** and selecting **Create New**. Add a URL pattern to be exempt and select **Exempt from Cache**.

You can also add URLs and addresses to be exempt from the CLI. Enter the following command to add www.example.com to the web cache exempt list.

```
config web-proxy url-match
   set cache-exemption enable
   set url-pattern www.example.com
end
```

Monitoring Web caching performance

The web cache monitor shows the percentage of web cache requests that retrieved content from the cache (hits) and the percentage that did not receive content from the cache (misses). A higher the number of hits usually indicates that the web cache is being more effective at reducing WAN traffic.

The web cache monitor also shows a graph of web traffic on the WAN and LAN. A lower WAN line on the graph indicates the web cache is reducing traffic on the WAN. The web cache monitor also displays the total number of web requests processed by the web cache.

To view the web cache monitor, go to **Monitor > Cache Monitor**.

Web cache monitor

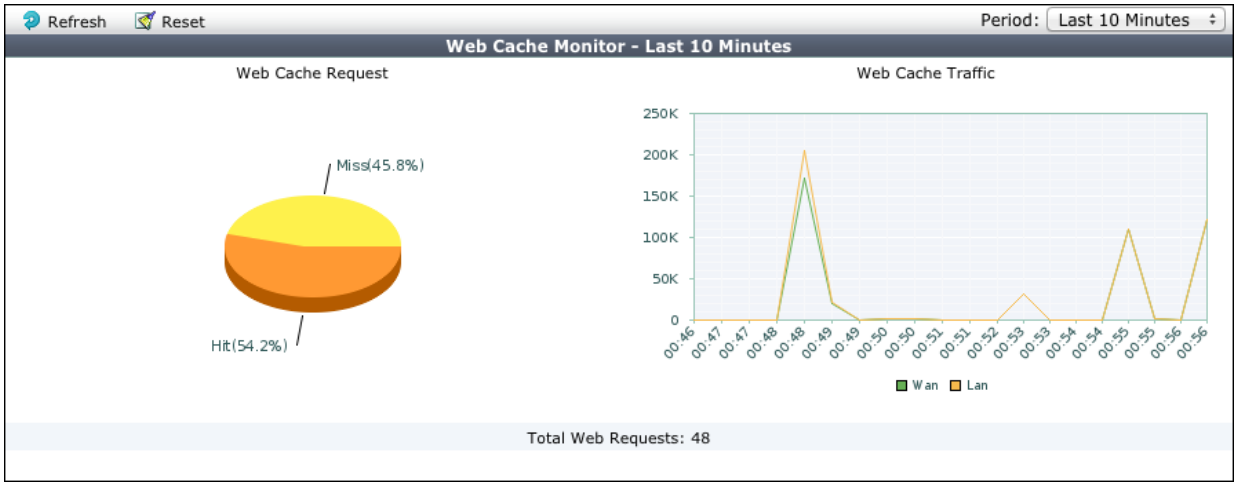

Example Web caching of HTTP and HTTPS Internet content for users on an internal network

This example describes how to configure web caching of HTTP and HTTPS for users on a private network connecting to the Internet.

Network topology and assumptions

This example includes a client network with subnet address 10.31.101.0 connecting to web servers on the Internet. All of the users on the private network access the Internet though a single general security policy on the FortiGate unit that accepts all sessions connecting to the Internet. Web caching for HTTP and HTTPS traffic is added to this security policy.

Since users on the private network have unrestricted access to the Internet and can be accessing many web servers the `webcache-https` is set to `any` and users may see error messages on their web browsers when accessing HTTPS content.

Initially, security profiles are not selected so the example caches all HTTP traffic on TCP port 80 and HTTPS traffic on port 443. The example also describes how to configure the security policy to cache HTTP traffic on port 80 and 8080 by adding a proxy options profile that looks for HTTP traffic on TCP ports 80 and 8080. The example also describes how to configure the security policy to cache HTTPS traffic on port 443 and 8443 using the same proxy options profile.

Example web caching topology

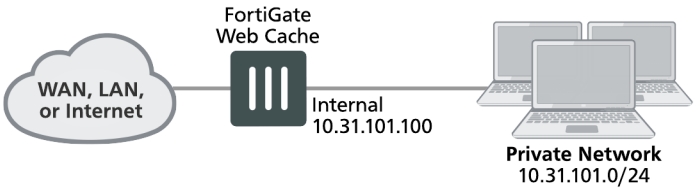

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Add HTTP web caching to the security policy that all users on the private network use to connect to the Internet.
2. Add HTTPS web caching.
3. Add a protocol options profile to look for HTTP traffic on ports 80 and 8080 and HTTPS traffic on ports 443 and 8443 and add this protocol options profile to the security policy.

If you perform any additional actions between procedures, your configuration may have different results.

Configuration Steps - web-based manager

Use the following steps to configure the example configuration from the FortiGate web-based manager.

To add HTTP web caching to a security policy

1. Go to **Policy & Objects > IPv4 Policy** and add a security policy that allows all users on the internal network to access the Internet.

Incoming Interface	Internal
Source Address	all
Outgoing Interface	wan1
Destination Address	all
Schedule	always
Service	ALL
Action	ACCEPT

2. Select **Enable NAT** and select **Use Destination Interface Address**.
3. Turn on **Web cache**.
4. Select **OK**.

To add HTTPS web caching

1. From the CLI enter the following command to add HTTPS web caching to the policy.

   Assume the index number of the policy is 5.

```
config firewall policy
   edit 5
      set webcache-https any
   end
```

To cache HTTP traffic on port 80 and 8080

1. Go to **Network > Explicit Proxy** and edit the **default** proxy options profile.

   You could also add a new profile.
2. Under **Protocol Port Mapping** enable **HTTP** and under **Inspection Ports** enter **80,8080**.
3. Go to **Policy & Objects > IPv4 Policy**, edit the security policy and

To cache HTTPS traffic on ports 443 and 8443

1. Go to **Security Profiles > SSL Inspection** and edit the **certificate-inspection** SSL/SSH inspection profile.

   You could also use the **deep-inspection** profile or add a new profile.
2. Under **SSL Inspection Options** select **Multiple Clients Connecting to Multiple Servers**.
3. Make sure **Inspect All Ports** is not selected.
4. Make sure **HTTPS** is turned on and enter **443,8443**.
5. From the CLI, enter the following command to add the **default** proxy options profile and the **certificate-inspection** SSL SSH profile to the firewall policy.

```
config firewall policy
   edit 5
      set utm-status enable
      set profile-protocol-options default
      set ssl-ssh-profile certificate-inspection
   end
```

You need to use the CLI to add the protocol options profile unless you also add a security profile that uses proxy-based inspection.

Configuration Steps - CLI

Use the following steps to configure the example configuration from the FortiGate CLI.

To add HTTP and HTTPS web caching to a security policy

1. Enter the following command to add a security policy that allows all users on the internal network to access the Internet and that includes web caching of HTTP and HTTPS traffic.

```
config firewall policy
   edit 0
      set srcintf internal
      set srcaddr all
      set dstintf wan1
      set distinf all
      set schedule always
      set service ANY
      set action accept
      set nat enable
```

```
    set webcache enable
    set webcache-https any
  end
```

To cache HTTP traffic on port 80 and 8080 and HTTPS traffic on ports 443 and 8443

1. Enter the following command to edit the **default** proxy options profile to configure it to look for HTTP traffic on ports 80 and 8080:

```
config firewall profile-protocol-options
  edit default
    config http
      set status enable
      set ports 80 8080
    end
```

2. Enter the following command to edit the **certification-inspection** SSL SSH options profile to configure it to look for HTTPS traffic on ports 443 and 8443:

```
config firewall ssl-ssh-profile
  edit certificate-inspection
    config https
      set status certificate-inspection
      set ports 443 8443
    end
```

3. Enter the following command to add the **default** proxy options profile and the **certificate-inspection** SSL SSH profile to the firewall policy.

```
config firewall policy
  edit 5
    set utm-status enable
    set profile-protocol-options default
    set ssl-ssh-profile certificate-inspection
  end
```

Example reverse proxy web caching and SSL offloading for an Internet web server using a static one-to-one virtual IP

This section describes configuring SSL offloading for a reverse proxy web caching configuration using a static one-to-one firewall virtual IP (VIP). While the static one-to-one configuration described in this example is valid, its also common to change the destination port of the unencrypted HTTPS traffic to a commonly used HTTP port such as 8080 using a port forwarding virtual IP.

Network topology and assumptions

In this configuration, clients on the Internet use HTTP and HTTPS to browse to a web server that is behind a FortiGate unit. A policy added to the FortiGate unit forwards the HTTP traffic to the web server. The policy also offloads HTTPS decryption and encryption from the web server so the web server only sees HTTP traffic.

The FortiGate unit also caches HTTP and HTTPS pages from the web server so when users access cached pages the web server does not see the traffic. Replies to HTTPS sessions are encrypted by the FortiGate unit before returning to the clients.

In this configuration, the FortiGate unit is operating as a web cache in reverse proxy mode. Reverse proxy caches can be placed directly in front of a web server. Web caching on the FortiGate unit reduces the number of requests that the web server must handle, therefore leaving it free to process new requests that it has not serviced before.

Using a reverse proxy configuration:

- avoids the capital expense of additional web servers by increasing the capacity of existing servers
- serves more requests for static content from web servers
- serves more requests for dynamic content from web servers
- reduces operating expenses including the cost of bandwidth required to serve content
- accelerates the response time of web servers and of page download times to end users.

When planning a reverse proxy implementation, the web server's content should be written so that it is “cache aware” to take full advantage of the reverse proxy cache.

In reverse proxy mode, the FortiGate unit functions more like a web server for clients on the Internet. Replicated content is delivered from the proxy cache to the external client without exposing the web server or the private network residing safely behind the firewall.

In this example, the site URL translates to IP address 192.168.10.1, which is the port2 IP address of the FortiGate unit. The port2 interface is connected to the Internet.

This example assumes that all HTTP traffic uses port 80 and all HTTPS traffic uses port 443.

The FortiGate unit includes the web server CA and an SSL server configuration for IP address 172.10.20.30 and port to 443. The name of the file containing the CA is Rev_Proxy_Cert_1.crt.

The destination address of incoming HTTP and HTTPS sessions is translated to the IP address of the web server using a static one-to-one virtual IP that performs destination address translation (DNAT) for the HTTP packets. The DNAT translates the destination address of the packets from 192.168.10.1 to 172.10.20.30 but does not change the destination port number.

When the SSL server on the FortiGate unit decrypts the HTTPS packets their destination port is changed to port 80.

Reverse proxy web caching and SSL offloading for an Internet web server using static one-to-one virtual IPs

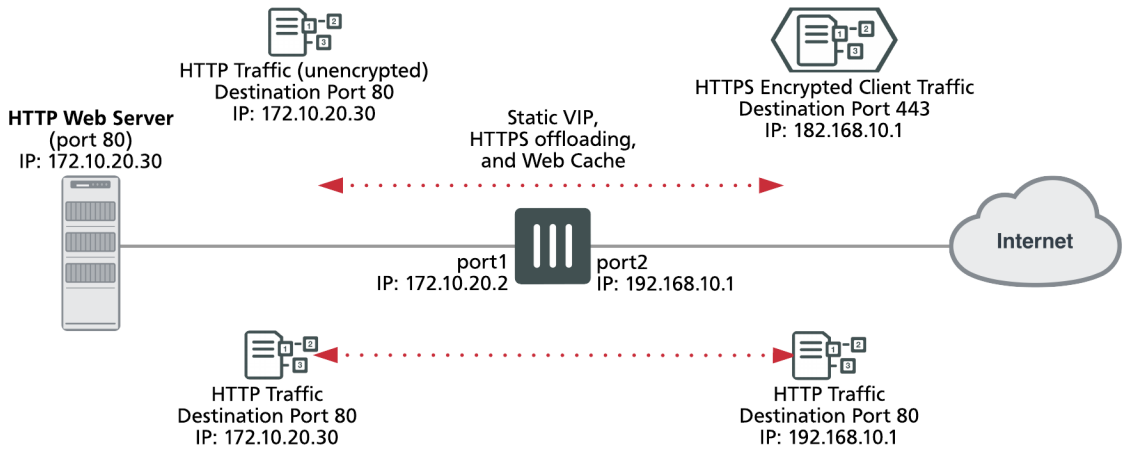

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Configure the FortiGate unit as a reverse proxy web cache server.
2. Configure the FortiGate unit for SSL offloading of HTTPS traffic.
3. Add an SSL server to offload SSL encryption and decryption for the web server.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

Configuration steps - web-based manager

To configure the FortiGate unit as a reverse proxy web cache server

1. Go to **Policy & Objects > Virtual IPs** and select **Create New** to add a static NAT virtual IP that translates destination IP addresses from 192.168.10.1 to 172.10.20.30 (and does not translate destination ports):

VIP Type	IPv4 VIP
Name	Reverse_proxy_VIP
Interface	port2
Type	Static NAT
Source Address Filter	Do not select.
External IP Address/Range	192.168.10.1
Mapped IP Address/Range	172.10.20.30
Port Forwarding	Do not select.

2. Select **OK**.
3. Go to **Policy & Objects > IPv4 Policy** and select **Create New** to add a port2 to port1 security policy that accepts HTTP and HTTPS traffic from the Internet.

   Do not select security profiles. Set the destination address to the virtual IP. You do not have to enable NAT.

Incoming Interface	port2
Source Address	all
Outgoing Interface	port1
Destination Address	Reverse_proxy_VIP
Schedule	always
Service	HTTP HTTPS
Action	ACCEPT

4. Turn on **Web Cache**.
5. Select **OK**.
6. From the CLI enter the following command to add HTTPS web caching to the security policy

   Assume the index number of the policy is 5.

```
config firewall policy
   edit 5
      set webcache-https ssl-server
   end
```

To configure the FortiGate unit to offload SSL encryption and cache HTTPS content

1. Go to **System > Certificates** and select **Import** to import the web server's CA.

   For **Type**, select **Local Certificate**. Select the **Browse** button to locate the file (example file name: Rev_Proxy_Cert_1.crt).

   The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.
2. Select **OK** to import the certificate.
3. From the CLI, enter the following command to add the SSL server and to add the server's certificate to the SSL server.

   The SSL server `ip` must match the destination address of the SSL traffic after being translated by the virtual IP (172.10.20.30) and the SSL server `port` must match the destination port of the SSL traffic (443). The SSL server operates in half mode since it performs a single-step conversion (HTTPS to HTTP or HTTP to HTTPS).

```
config firewall ssl-server
   edit rev_proxy_server
      set ip 172.10.20.30
      set port 443
      set ssl-mode half
      set ssl-cert Rev_Proxy_Cert_1
```

```
end
```

Configuration steps - CLI

To configure the FortiGate unit as a reverse proxy web cache server

1. Enter the following command to add a static NAT virtual IP that translates destination IP addresses from 192.168.10.1 to 172.10.20.30 (and does not translate destination ports):

```
config firewall vip
   edit Reverse_proxy_VIP
      set extintf port2
      set type static-nat
      set extip 192.168.10.1
      set mappedip 172.10.20.30
   end
```

2. Enter the following command to add a port2 to port1 security policy that accepts HTTP and HTTPS traffic from the Internet. Enable web caching and HTTPS web caching.

   Do not select security profiles. Set the destination address to the virtual IP. You do not have to enable NAT.

```
config firewall policy
   edit 0
      set srcintf port2
      set srcaddr all
      set dstintf port1
      set dstaddr Reverse_proxy_VIP
      set schedule always
      set service HTTP HTTPS
      set action accept
      set webcache enable
      set webcache-https ssl-server
   end
```

To add an SSL server to offload SSL encryption and decryption for the web server

1. Place a copy of the web server's CA (file name Rev_Proxy_Cert_1.crt) in the root folder of a TFTP server.
2. Enter the following command to import the web server's CA from a TFTP server. The IP address of the TFTP server is 10.31.101.30:

```
execute vpn certificate local import tftp Rev_Proxy_Cert_1.crt 10.31.101.30
```

   The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.

3. From the CLI, enter the following command to add the SSL server.

   The SSL server `ip` must match the destination address of the SSL traffic after being translated by the virtual IP (172.10.20.30) and the SSL server `port` must match the destination port of the SSL traffic (443). The SSL server operates in half mode since it performs a single-step conversion (HTTPS to HTTP or HTTP to HTTPS).

```
config firewall ssl-server
   edit rev_proxy_server
      set ip 172.10.20.30
      set port 443
      set ssl-mode half
      set ssl-cert Rev_Proxy_Cert_1
   end
```

4. Configure other `ssl-server` settings that you may require for your configuration.

FortiClient WAN optimization

FortiClient WAN optimization supports protocol optimization and byte caching in IPsec VPN and SSL VPN tunnels between FortiClient and a FortiGate unit. To add WAN optimization to FortiClient, configure FortiClient Advanced settings and enable WAN optimization. This setting can then apply WAN optimization to any IPsec or SSL VPN tunnel between FortiClient and FortiGate, if the FortiGate IPsec or SSL VPN configuration also includes WAN optimization.

When FortiClient with WAN optimization enabled attempts to connect a server-side FortiGate unit, FortiClient automatically detects if WAN optimization has been added to the FortiGate tunnel configuration. If WAN optimization is detected and FortiClient can successfully negotiate with the FortiGate unit, WAN optimization starts.

FortiClient WAN optimization topology

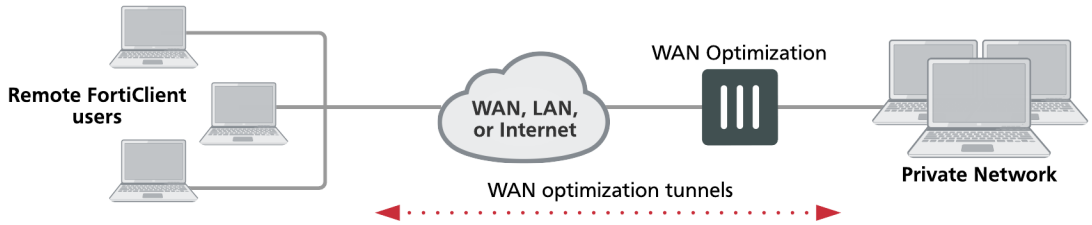

FortiClient WAN optimization over IPsec VPN configuration example

This example shows how to add WAN optimization to a FortiClient IPsec VPN. The IPsec VPN tunnel allows remote FortiClient users to connect to the internal network behind the FortiGate unit.

Example FortiClient WAN optimization configuration

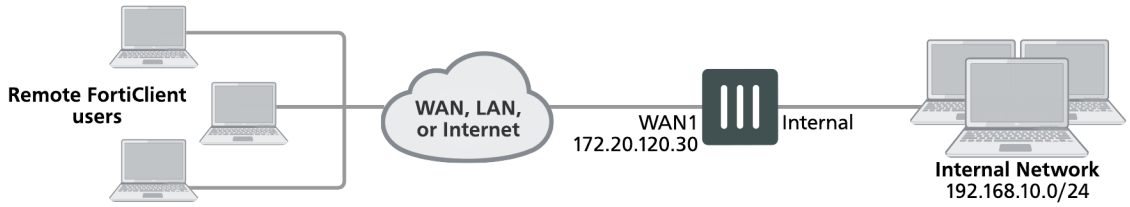

To configure the FortiGate unit

Because computers running FortiClient can have IP addresses that change often, it is usually not practical to add FortiClient peers to the FortiGate WAN optimization peer list. Instead, a FortiGate unit that accepts WAN optimization tunnel requests from FortiClient is usually configured to accept any peer. This example does this by adding a WAN optimization authentication group with **Peer acceptance** set to **Accept Any Peer**.

In addition this example includes a **wanopt** to **internal** policy to allow WAN optimization traffic reach the internal network. Finally passive WAN optimization is added to the ssl.root policy because WAN optimization is accepting traffic from the IPsec VPN tunnel.

1. Go to **WAN Opt. & Cache > Authentication Groups** and select **Create New**.
2. Configure the WAN optimization authentication group:

Name	auth-fc
Authentication Method	Certificate
Certificate	Fortinet_Firmware
Peer Acceptance	Accept Any Peer

3. Select **OK**.
4. Go to **WAN Opt. & Cache > Profiles** and select **Create New** (select the + button).
5. Add a profile for FortiClient WAN optimization sessions:

Name	Fclient_Pro
Transparent Mode	Select
Authentication Group	auth-fc

6. Select any Protocols and any settings for each protocol.
7. Select **OK**.
8. Go to **Policy & Objects > Addresses** and select **Create New** to add a firewall address for the internal network that FortiClient users can access.

Category	Address
Address Name	Internal-Server-Net
Type	IP Range
Subnet / IP Range	192.168.10.0/24
Interface	internal

9. Enter the following CLI command to add an explicit proxy policy to accept WAN optimization tunnel connections.

```
configure firewall explicit-proxy-policy
   edit 0
      set proxy wanopt
      set dstintf internal
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service ALL
   next
end
```

To set up IPsec VPN to support WAN optimization

1. Go to **VPN > IPsec Wizard**, enter a **Name** for the IPsec VPN and select **Dialup - FortiClient (Windows, Mac OS, Android)**.
2. Follow the wizard steps to configure the VPN. No special WAN optimization settings are required.
3. Go to **Policy & Objects > IPv4 Policy** and edit the policy created by the wizard.

   This policy has the IPsec VPN interface created by the wizard as the source interface.
4. Turn on **WAN Optimization** and configure the following settings:

Enable WAN Optimization	passive
Passive Option	default

5. Select **OK**.

To configure FortiClient and start the WAN optimization SSL VPN connection

1. Open FortiClient, configure **Advanced** settings, and select **Enable WAN optimization**.
2. Add a new IPsec VPN connection.

   Set the Server to the WAN1 IP address of the FortiGate unit (172.20.120.30 in this example).

   No other settings are required for this example. You can add authentication in the form of a user name and password if required by the FortiGate unit.
3. Start the IPsec VPN tunnel.

   You should be connected to the IPsec VPN tunnel and traffic in it should be optimized.

The FortiGate explicit web proxy

You can use the FortiGate explicit web proxy to enable explicit proxying of IPv4 and IPv6 HTTP, and HTTPS traffic on one or more FortiGate interfaces. The explicit web proxy also supports proxying FTP sessions from a web browser and proxy auto-config (PAC) to provide automatic proxy configurations for explicit web proxy users. From the CLI you can also configure the explicit web proxy to support SOCKS sessions from a web browser.

The explicit web and FTP proxies can be operating at the same time on the same or on different FortiGate interfaces.

If explicit web proxy options are not visible on the web-based manager, go to **System > Feature Select** and turn on **Explicit Proxy**.

In most cases you would configure the explicit web proxy for users on a network by enabling the explicit web proxy on the FortiGate interface connected to that network. Users on the network would configure their web browsers to use a proxy server for HTTP and HTTPS, FTP, or SOCKS and set the proxy server IP address to the IP address of the FortiGate interface connected to their network. Users could also enter the PAC URL into their web browser PAC configuration to automate their web proxy configuration using a PAC file stored on the FortiGate unit.

Enabling the explicit web proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address.

If the FortiGate unit is operating in Transparent mode, users would configure their browsers to use a proxy server with the FortiGate management IP address.

If the FortiGate unit is operating with multiple VDOMs the explicit web proxy is configured for each VDOM.

The web proxy receives web browser sessions to be proxied at FortiGate interfaces with the explicit web proxy enabled. The web proxy uses FortiGate routing to route sessions through the FortiGate unit to a destination interface. Before a session leaves the exiting interface, the explicit web proxy changes the source addresses of the session packets to the IP address of the exiting interface. When the FortiGate unit is operating in Transparent mode the explicit web proxy changes the source addresses to the management IP address. You can configure the explicit web proxy to keep the original client IP address. See Preventing the explicit web proxy from changing source addresses on page 113.

For more information about explicit web proxy sessions, see Explicit web proxy sessions and user limits on page 118.

Example explicit web proxy topology

To allow all explicit web proxy traffic to pass through the FortiGate unit you can set the explicit web proxy default firewall policy action to accept. However, in most cases you would want to use security policies to control explicit web proxy traffic and apply security features such as access control/authentication, virus scanning, web filtering, application control, and traffic logging. You can do this by keeping the default explicit web proxy security policy action to deny and then adding web-proxy security policies.

You can also change the explicit web proxy default security policy action to accept and add explicit web proxy security policies. If you do this, sessions that match web-proxy security policies are processed according to the security policy settings. Connections to the explicit web proxy that do not match a web-proxy security policy are allowed with no restrictions or additional security processing. This configuration is not recommended and is not a best practice.

The explicit web-proxy can accept VIP addresses for destination address. If an external IP matches a VIP policy, the IP is changed to the mapped-IP of the VIP.

Web-proxy policies can selectively accept or deny traffic, apply authentication, enable traffic logging, and use security profiles to apply virus scanning, web filtering, IPS, application control, DLP, and SSL/SSH inspection to explicit web proxy traffic.

You cannot configure IPsec, SSL VPN, or Traffic shaping for explicit web proxy traffic. Web Proxy policies can only include firewall addresses not assigned to a FortiGate unit interface or with interface set to **Any**. (On the web-based manager you must set the interface to **Any**. In the CLI you must `unset` the `associated-interface`.)

Authentication of explicit web proxy sessions uses HTTP authentication and can be based on the user's source IP address or on cookies from the user's web browser. For more information, see Explicit web proxy authentication on page 107.

To use the explicit web proxy, users must add the IP address of a FortiGate interface on which the explicit web proxy is enabled and the explicit web proxy port number (default 8080) to the proxy configuration settings of their web browsers.

On FortiGate units that support it, you can also enable web caching for explicit web proxy sessions.

For the time being, traffic shaping is not supported per policy for explicit proxy. For explicit proxy traffic, traffic shaping can be carried out per interface.

General explicit web proxy configuration steps

You can use the following general steps to configure the explicit web proxy.

To enable the explicit web proxy - web-based manager:

1. Go to **Network > Explicit Proxy** and enable **Explicit Web Proxy**. From here you can optionally change the HTTP port that the proxy listens on (the default is 8080) and optionally specify different ports for HTTPS, FTP, PAC, and other options.
2. Optionally enable **IPv6 Explicit Proxy** to turn on the explicit web proxy for IPv6 traffic.

If you enable both the IPv4 and the IPv6 explicit web proxy you can combine IPv4 and IPv6 addresses in a single explicit web proxy policy to allow both IPv4 and IPv6 traffic through the proxy.

3. Select **Apply**.

4. Go to **Network > Interfaces** and select one or more interfaces for which to enable the explicit web proxy. Edit the interface and select **Enable Explicit Web Proxy**.

Enabling the explicit web proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address. If you enable the proxy on such an interface make sure authentication is required to use the proxy.

5. Go to **Policy & Objects > Addresses** and select **Create New** to add a firewall address that matches the source address of packets to be accepted by the explicit proxy.

Category	Address
Name	Internal_subnet
Type	IP Range
Subnet / IP Range	10.31.101.1 - 10.31.101.255
Interface	any*

*The **Interface** must be set to **Any**.

You can also set the **Type** to **URL Pattern (Explicit Proxy)** to add a destination URL that is only used by the explicit proxy. For example, to create an explicit policy that only allows access to Fortinet.com:

Category	Address
Name	Fortinet-web-sites
Type	URL Pattern (Explicit Proxy)
URL Pattern	fortinet.com
Interface	any

6. Go to **Policy & Objects > Explicit Proxy Policy** and select **Create New.** Configure the policy as required to accept the traffic that you want to be allowed to use the explicit web proxy.

The source address of the policy must match the client's source IP addresses. The interface of this firewall address must be set to **any**.

The destination address of the policy should match the IP addresses of web sites that clients are connecting to. Usually the destination address would be **all** if proxying Internet web browsing. You could also specify a URL

firewall address to limit the policy to allowing access to this URL.

If **Default Firewall Policy Action** is set to **Deny** (under **Network > Explicit Proxy**), traffic sent to the explicit web proxy that is not accepted by a web-proxy policy is dropped. If **Default Firewall Policy Action** is set to **Allow** then all web-proxy sessions that don't match with a security policy are allowed.

For example, the following security policy allows users on an internal network to access fortinet.com websites through the wan1 interface of a FortiGate unit.

Explicit Proxy Type	Web
Source Address	Internal_subnet
Outgoing Interface	wan1
Destination Address	Fortinet-web-sites
Schedule	always
Action	ACCEPT

Add security profiles as required.

7. Select Create New to add another explicit web proxy and set the **Action** to **AUTHENTICATE** to require authentication to access the explicit web proxy. For example:

Explicit Proxy Type	Web
Source Address	Internal_subnet
Outgoing Interface	wan1
Destination Address	Fortinet-web-sites
Schedule	always
Action	AUTHENTICATE

Select **Create New** to add an **Authentication Rule** and configure the rule as follows:

Groups	Proxy-Group
Source User(s)	(optional)
Schedule	always

Add security profiles as required and select **OK**.

You can add multiple user identity policies to apply different authentication for different user groups and users and also apply different UTM and logging settings for different user groups.

You can change the **User Authentication Options** if required. In most cases you can accept the defaults.

8. Optionally enable Web Caching.
9. Select **OK**.

To enable the explicit web proxy - CLI:

1. Enter the following command to turn on the IPv4 and IPv6 explicit web proxy for HTTP and HTTPS traffic.

```
config web-proxy explicit
   set status enable
   set ipv6-status enable
end
```

   You can also enter the following command to enable the web proxy for FTP sessions in a web browser.

```
config web-proxy explicit
   set ftp-over-http enable
end
```

   The default explicit web proxy configuration has `sec-default-action` set to `deny` and requires you to add a security policy to allow access to the explicit web proxy.

2. Enter the following command to enable the explicit web proxy for the internal interface.

```
config system interface
   edit internal
      set explicit-web-proxy enable
   end
end
```

3. Use the following command to add a firewall address that matches the source address of users who connect to the explicit web proxy.

```
config firewall address
   edit Internal_subnet
      set type iprange
      set start-ip 10.31.101.1
      set end-ip 10.31.101.255
   end
```

   The source address for a web-proxy security policy cannot be assigned to a FortiGate interface.

4. Optionally use the following command to add a destination URL that is only used by the explicit proxy. For example, to create an explicit policy that only allows access to Fortinet.com:

```
config firewall address
   edit Fortinet-web-sites
      set type url
      set url fortinet.com
   end
```

5. Use the following command to add an explicit web proxy policy that allows all users on the internal subnet to use the explicit web proxy for connections through the wan1 interface to the Internet.

```
config firewall explicit-proxy-policy
   edit 0
      set proxy web
      set dstintf wan1
      set scraddr Internal_subnet
      set dstaddr all
```

```
        set action accept
        set service webproxy
        set schedule always
    end
```

6. Use the following command to add an explicit web proxy policy that allows authenticated users on the internal subnet to use the explicit web proxy for connections through the wan1 interface to the Internet.

```
config firewall explicit-proxy-policy
   edit 0
      set proxy web
      set dstintf wan1
      set scraddr Internal_subnet
      set dstaddr Fortinet-web-sites
      set action accept
      set service webproxy
      set schedule always
      set identity-based enable
         config identity-based-policy
            edit 1
               set groups Proxy-group
               set schedule always
      end
end
```

7. Use the following command to change global web proxy settings, for example to set the maximum request length for the explicit web proxy to 10:

```
config web-proxy global
   set max-request-length 10
end
```

Explicit proxy firewall address types

Explicit proxy firewall address types improve granularity over header matching for explicit web proxy policies. You can enable this option using the **Show in Address List** button on the Address and Address Group New/Edit forms under **Policy & Objects > Addresses**.

The following address types are available:

- **URL Pattern** - destination address
- **Host Regex Match** - destination address
- **URL Category** - destination address (URL filtering)
- **HTTP Method** - source address
- **User Agent** - source address
- **HTTP Header** - source address
- **Advanced (Source)** - source address (combines User Agent, HTTP Method, and HTTP Header)
- **Advanced (Destination)** - destination address (combines Host Regex Match and URL Category)

Proxy auto-config (PAC) configuration

A proxy auto-config (PAC) file defines how web browsers can choose a proxy server for receiving HTTP content. PAC files include the FindProxyForURL(url, host) JavaScript function that returns a string with one or more access method specifications. These specifications cause the web browser to use a particular proxy server or to connect directly.

To configure PAC for explicit web proxy users, you can use the port that PAC traffic from client web browsers use to connect to the explicit web proxy. explicit web proxy users must configure their web browser's PAC proxy settings to use the PAC port.

PAC File Content

You can edit the default PAC file from the web-based manager or use the following command to upload a custom PAC file:

```
config web-proxy explicit
   set pac-file-server-status enable
   set pac-file-data <pac_file_str>
end
```

Where `<pac_file_str>` is the contents of the PAC file. Enter the PAC file text in quotes. You can copy the contents of a PAC text file and paste the contents into the CLI using this option. Enter the command followed by two sets of quotes then place the cursor between the quotes and paste the file content.

The maximum PAC file size is 256 kbytes. If your FortiGate unit is operating with multiple VDOMs each VDOM has its own PAC file. The total amount of FortiGate memory available to store all of these PAC files 2 MBytes. If this limit is reached you will not be able to load any additional PAC files.

You can use any PAC file syntax that is supported by your users's browsers. The FortiGate unit does not parse the PAC file.

To use PAC, users must add an automatic proxy configuration URL (or PAC URL) to their web browser proxy configuration. The default FortiGate PAC file URL is:

```
http://<interface_ip>:<PAC_port_int>/<pac_file_str>
```

For example, if the interface with the explicit web proxy has IP address 172.20.120.122, the PAC port is the same as the default HTTP explicit web proxy port (8080) and the PAC file name is proxy.pac the PAC file URL would be:

```
http://172.20.120.122:8080/proxy.pac
```

From the CLI you can use the following command to display the PAC file URLs:

```
get web-proxy explicit
```

Unknown HTTP version

You can select the action to take when the proxy server must handle an unknown HTTP version request or message. Set unknown HTTP version to Reject or Best Effort. Best Effort attempts to handle the HTTP traffic as best as it can. Reject treats known HTTP traffic as malformed and drops it. The Reject option is more secure.

Authentication realm

You can enter an authentication realm to identify the explicit web proxy. The realm can be any text string of up to 63 characters. If the realm includes spaces enclose it in quotes. When a user authenticates with the explicit web proxy the HTTP authentication dialog includes the realm so you can use the realm to identify the explicitly web proxy for your users.

Implementing Botnet features

The option `scan-botnet-connections` can be added to an explicit proxy policy.

CLI Syntax:

```
config firewall explicit-proxy-policy
   edit <policy_id>
      set scan-botnet-connections [disable|block|monitor]
end
```

where:

- `disable` means do not scan connections to botnet servers.
- `block` means block connections to botnet servers.
- `monitor` means log connections to botnet servers.

Other explicit web proxy options

You can change the following explicit web proxy options as required by your configuration.

HTTP port, HTTPS port, FTP port, PAC port	The TCP port that web browsers use to connect to the explicit proxy for HTTP, HTTPS, FTP and PAC services. The default port is 8080 for all services. By default HTTPS, FTP. and PAC use the same port as HTTP. You can change any of these ports as required. Users configuring their web browsers to use the explicit web proxy should add the same port numbers to their browser configurations.
Proxy FQDN	Enter the fully qualified domain name (FQDN) for the proxy server. This is the domain name to enter into browsers to access the proxy server.
Max HTTP request length	Enter the maximum length of an HTTP request in Kbytes. Larger requests will be rejected.
Max HTTP message length	Enter the maximum length of an HTTP message in Kbytes. Larger messages will be rejected.

Configuring an external IP address for the IPv4 explicit web proxy

You can use the following command to set an external IP address (or pool) that will be used by the explicit web proxy policy.

```
config web-proxy explicit
   set status enable
   set outgoing-ip <ip1> <ip2> ... <ipN>
end
```

Configuring an external IP address for the IPv6 explicit web proxy

You can use the following command to set an external IP address (or pool) that will be used by the explicit web proxy policy.

```
config web-proxy explicit
   set status enable
   set outgoing-ipv6 <ip1> <ip2> ... <ipN>
end
```

Restricting the IP address of the IPv4 explicit web proxy

You can use the following command to restrict access to the explicit web proxy using only one IP address. The IP address that you specify must be the IP address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit FTP proxy is enabled on an interface with multiple IP addresses.

For example, to require uses to connect to the IP address 10.31.101.100 to connect to the explicit HTTP proxy:

```
config web-proxy explicit
   set incoming-ip 10.31.101.100
end
```

Restricting the outgoing source IP address of the IPv4 explicit web proxy

You can use the following command to restrict the source address of outgoing web proxy packets to a single IP address. The IP address that you specify must be the IP address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit HTTP proxy is enabled on an interface with multiple IP addresses.

For example, to restrict the outgoing packet source address to 172.20.120.100:

```
config http-proxy explicit
   set outgoing-ip 172.20.120.100
end
```

Restricting the IP address of the explicit IPv6 web proxy

You can use the following command to restrict access to the IPv6 explicit web proxy to use only one IP6 IP address. The IPv6 address that you specify must be the IPv6 address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit web proxy is enabled on an interface with multiple IPv6 addresses.

For example, to require uses to connect to the IPv6 address 2001:db8:0:2::30 to connect to the explicit IPv6 HTTP proxy:

```
config web-proxy explicit
   set incoming-ipv6 2001:db8:0:2::30
end
```

Restricting the outgoing source IP address of the IPv6 explicit web proxy

You can use the following command to restrict the source address of outgoing web proxy packets to a single IPv6 address. The IP address that you specify must be the IPv6 address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit HTTP proxy is enabled on an interface with multiple IPv6 addresses.

For example, to restrict the outgoing packet source address to 2001:db8:0:2::50:

```
config http-proxy explicit
   set outgoing-ipv6 2001:db8:0:2::50
end
```

Proxy chaining (web proxy forwarding servers)

For the explicit web proxy you can configure web proxy forwarding servers to use proxy chaining to redirect web proxy sessions to other proxy servers. Proxy chaining can be used to forward web proxy sessions from the FortiGate unit to one or more other proxy servers on your network or on a remote network. You can use proxy chaining to integrate the FortiGate explicit web proxy with an web proxy solution that you already have in place.

A FortiGate unit can forward sessions to most web proxy servers including a remote FortiGate unit with the explicit web proxy enabled. No special configuration of the explicit web proxy on the remote FortiGate unit is required.

You can deploy the explicit web proxy with proxy chaining in an enterprise environment consisting of small satellite offices and a main office. If each office has a FortiGate unit, users at each of the satellite offices can use their local FortiGate unit as an explicit web proxy server. The satellite office FortiGate units can forward explicit web proxy sessions to an explicit web proxy server at the central office. From here the sessions can connect to web servers on the Internet.

FortiGate proxy chaining does not support authenticating with the remote forwarding server.

Adding a web proxy forwarding server

To add a forwarding server, select **Create New** in the **Web Proxy Forwarding Servers** section of the **Explicit Proxy** page by going to **Network > Explicit Proxy**.

Server Name	Enter the name of the forwarding server.
Proxy Address	Enter the IP address of the forwarding server.
Proxy Address Type	Select the type of IP address of the forwarding server. A forwarding server can have an FQDN or IP address.
Port	Enter the port number on which the proxy receives connections. Traffic leaving the FortiGate explicit web proxy for this server has its destination port number changed to this number.
Server Down action	Select what action the explicit web proxy to take if the forwarding server is down. **Block** means if the remote server is down block traffic. **Use Original Server** means do not forward traffic to the forwarding sever but instead forward it from the FortiGate to its destination. In other words operate as if there is no forwarding server configured.
Enable Health Monitor **Health Check Monitor Site**	Select to enable health check monitoring and enter the address of a remote site. See "Web proxy forwarding server monitoring and health checking".

Use the following CLI command to add a web proxy forwarding server named `fwd-srv` at address `proxy.example.com` and port 8080.

```
config web-proxy forward-server
   edit fwd-srv
      set addr-type fqdn
      set fqdn proxy.example.com
      set port 8080
   end
```

Web proxy forwarding server monitoring and health checking

By default, a FortiGate unit monitors web proxy forwarding server by forwarding a connection to the remote server every 10 seconds. If the remote server does not respond it is assumed to be down. Checking continues and when the server does send a response the server is assumed to be back up. If you configure health checking, every 10 seconds the FortiGate unit attempts to get a response from a web server by connecting through the remote forwarding server.

You can configure health checking for each remote server and specify a different website to check for each one.

If the remote server is found to be down you can configure the FortiGate unit to block sessions until the server comes back up or to allow sessions to connect to their destination, bypassing the remote forwarding server. You cannot configure the FortiGate unit to fail over to another remote forwarding server.

Configure the server down action and enable health monitoring from the web-based manager by going to **Network > Explicit Proxy**, selecting a forwarding server, and changing the server down action and changing the health monitor settings.

Use the following CLI command to enable health checking for a web proxy forwarding server and set the server down option to bypass the forwarding server if it is down.

```
config web-proxy forward-server
   edit fwd-srv
      set healthcheck enable
      set monitor http://example.com
      set server-down-option pass
   end
```

Grouping forwarding servers and load balancing traffic to them

You can add multiple web proxy forwarding servers to a forwarding server group and then add the server group to an explicit web proxy policy instead of adding a single server. Forwarding server groups are created from the FortiGate CLI but can be added to policies from the web-based manager (or from the CLI).

When you create a forwarding server group you can select a load balancing method to control how sessions are load balanced to the forwarding servers in the server group. Two load balancing methods are available:

- **Weighted** load balancing sends more sessions to the servers with higher weights. You can configure the weight for each server when you add it to the group.
- **Least-session** load balancing sends new sessions to the forwarding server that is processing the fewest sessions.

When you create a forwarding server group you can also enable **affinity**. Enable affinity to have requests from the same client processed by the same server. This can reduce delays caused by using multiple servers for a single multi-step client operation. Affinity takes precedence over load balancing.

You can also configure the behavior of the group if all of the servers in the group are down. You can select to **block** traffic or you can select to have the traffic **pass** through the FortiGate explicit proxy directly to its destination instead of being sent to one of the forwarding servers.

Use the following command to add a forwarding server group that users weighted load balancing to load balance traffic to three forwarding servers. Server weights are configured to send most traffic to server2. The group has affinity enabled and blocks traffic if all of the forward servers are down:

```
config web-proxy forward-server
   edit server_1
      set ip 172.20.120.12
      set port 8080
   next
   edit server_2
      set ip 172.20.120.13
      set port 8000
   next
   edit server_3
      set ip 172.20.120.14
      set port 8090
   next
   end
```

```
config web-proxy forward-server-group
   edit New-fwd-group
      set affinity enable
      set ldb-method weight
      set group-down-option block
         config server-list
            edit server_1
               set weight 10
            next
            edit server_2
               set weight 40
            next
            edit server_3
               set weight 10
            next
   end
```

Adding proxy chaining to an explicit web proxy policy

You enable proxy chaining for web proxy sessions by adding a web proxy forwarding server or server group to an explicit web proxy policy. In a policy you can select one web proxy forwarding server or server group. All explicit web proxy traffic accepted by this security policy is forwarded to the specified web proxy forwarding server or server group.

To add an explicit web proxy forwarding server - web-based manager:

1. Go to **Policy & Objects > Explicit Proxy Policy** and select **Create New**.
2. Configure the policy:

Explicit Proxy Type	Web
Source Address	Internal_subnet
Outgoing Interface	wan1
Destination Address	all
Schedule	always
Action	ACCEPT
Web Proxy Forwarding Server	Select, fwd-srv

3. Select **OK** to save the security policy.

To add an explicit web proxy forwarding server - CLI:

1. Use the following command to add a security policy that allows all users on the 10.31.101.0 subnet to use the explicit web proxy for connections through the wan1 interface to the Internet. The policy forwards web proxy sessions to a remote forwarding server named `fwd-srv`

```
config firewall explicit-proxy-policy
   edit 0
      set proxy web
```

```
    set dstintf wan1
    set scraddr Internal_subnet
    set dstaddr all
    set action accept
    set schedule always
    set webproxy-forward-server fwd-srv
end
```

Adding disclaimer messages to explicit proxy policies

This feature allows you to create user exceptions for specific URL categories (including warning messages) based on user groups. The **Disclaimer Options** are configured under **Policy & Objects > Explicit Proxy Policy**.

You can also configure a disclaimer for each Authentication Rule by setting **Action** to **Authenticate**.

Disclaimer explanations

- **Disable:** No disclaimer (default setting).
- **By Domain:** The disclaimer will be displayed on different domains. The explicit web proxy will check the referring header to mitigate the javascript/css/images/video/etc page.
- **By Policy:** The disclaimer will be displayed ifa the HTTP request matches a different explicit firewall policy.
- **By User:** The disclaimer will be displayed when a new user logs on.

Explicit web proxy authentication

You can add authentication to explicit web proxy policies to control access to the explicit web proxy and to identify users and apply different UTM features to different users.

Authentication of web proxy sessions uses HTTP basic and digest authentication as described in RFC 2617 (HTTP Authentication: Basic and Digest Access Authentication) and prompts the user for credentials from the browser allowing individual users to be identified by their web browser instead of IP address. HTTP authentication allows the FortiGate unit to distinguish between multiple users accessing services from a shared IP address.

You can also select IP-based authentication to authenticate users according to their source IP address in the same way as normal firewall policies.

IP-Based authentication

IP-based authentication applies authentication by source IP address. For the explicit web proxy, IP authentication is compatible with basic, digest, NTLM, FSSO, or RSSO authentication methods. Once a user authenticates, all sessions to the explicit web proxy from that user's IP address are assumed to be from that user and are accepted until the authentication timeout ends or the session times out.

This method of authentication is similar to standard (non-web proxy) firewall authentication and may not produce the desired results if multiple users share IP addresses (such as in a network that uses virtualization solutions or includes a NAT device between the users and the explicit web proxy).

To configure IP-based authentication, add an explicit web proxy security policy, set the **Action to AUTHENTICATION**, and select **Enable IP Based Authentication** is selected.

Use the following CLI command to add IP-based authentication to a web proxy security policy. IP-based authentication is selected by setting `ip-based` to `enable`.

```
config firewall explicit-proxy-policy
   edit 0
      set proxy web
      set scraddr User_network
      set dstintf port1
      set dstaddr all
      set action accept
      set identity-based enable
      set ip-based enable
         config identity-based-policy
            edit 0
               set groups Internal_users
               set users dwhite rlee
               set schedule always
            end
```

Per session authentication

If you don't select **IP Based** the explicit web proxy applies HTTP authentication per session. This authentication is browser-based. When a user enters a user name and password in their browser to authenticate with the explicit web proxy, this information is stored by the browser in a session cookie. Each new session started by the same web browser uses the session cookie for authentication. When the session cookie expires the user has to re-authenticate. If the user starts another browser on the same PC or closes and then re-opens their browser they have to authenticate again.

Since the authentication is browser-based, multiple clients with the same IP address can authenticate with the proxy using their own credentials. HTTP authentication provides authentication for multiple user sessions from the same source IP address. This can happen if there is a NAT device between the users and the FortiGate unit. HTTP authentication also supports authentication for other configurations that share one IP address among multiple users. These includes Citrix products and Windows Terminal Server and other similar virtualization solutions.

To configure per session authentication, add a explicit web proxy policy, set the **Action** to **AUTHENTICATE**, and make sure **Enable IP Based Authentication** is not selected.

Use the following CLI command to add per session authentication to a security policy. Per session authentication is selected by setting `ip-based` to `disable`.

```
config firewall explicit-proxy-policy
   edit 0
      set proxy web
      set scraddr User_network
      set dstintf port1
      set dstaddr all
      set action accept
      set identity-based enable
      set ip-based disable
         config identity-based-policy
            edit 0
               set groups Internal_users
               set users dwhite rlee
               set schedule always
            end
```

```
end
```

Per session HTTP authentication

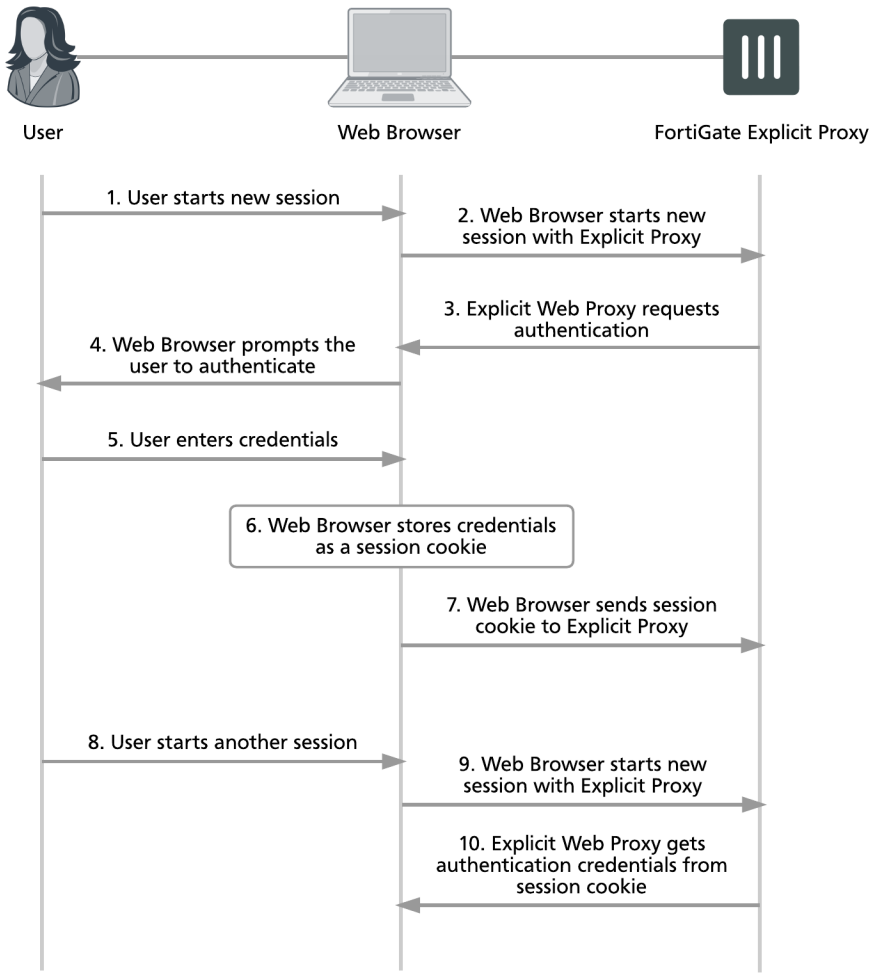

Transaction-based authentication

Multiple HTTP transactions (request/response) may be pipelined in the same TCP connection. Typically, all HTTP transactions of a TCP connection are considered as belonging to the same user. However, some devices (e.g., load balancers) may send HTTP transactions of different users to the same TCP connection and to explicit proxy. In order to support this deployment case, transaction-based authentication can be implemented to require each HTTP transaction to be authenticated separately.

To implement transaction-based authentication in the CLI:

```
config firewall explicit-proxy-policy
   edit <id>
      set transaction-based enable
   next
end
```

Security profiles, threat weight, device identification, and the explicit web proxy

You can apply all security profiles to explicit web proxy sessions. This includes antivirus, web filtering, intrusion protection (IPS), application control, data leak prevention (DLP), and SSL/SSH inspection. Security profiles are applied by selecting them in an explicit web proxy policy or in authentication rules added to web proxy policies.

Traffic accepted by explicit web proxy policies contributes to threat weight data.

The explicit web proxy is not compatible with device identification.

Since the traffic accepted by the explicit web proxy is known to be either HTTP, HTTPS, or FTP over HTTP and since the ports are already known by the proxy, the explicit web proxy does not use all of the SSL/SSH inspection options. The explicit web proxy does support the following proxy options:

- Enable chunked bypass
- HTTP oversized file action and threshold

The explicit web proxy does not support the following proxy options:

- Client comforting
- Server comforting
- Monitor content information from dashboard. URLs visited by explicit web proxy users are not added to dashboard usage and log and archive statistics widgets.

For explicit web proxy sessions, the FortiGate unit applies antivirus scanning to HTTP POST requests and HTTP responses. The FortiGate unit starts virus scanning a file in an HTTP session when it receives a file in the body of an HTML request. The explicit web proxy can receive HTTP responses from either the originating web server or the FortiGate web cache module.

Web Proxy firewall services and service groups

Configure web proxy services by selecting **Explicit Proxy** when configuring a service. Web proxy services can be selected in a explicit web proxy policy when adding one from the CLI. If you add a policy from the web-based manager the service is set to the **webproxy** service. The webproxy service should be used in most cases, it matches with any traffic with any port number. However, if you have special requirements, such as using a custom protocol type or a reduced port range or need to add an IP/FQDN to an explicit proxy service you can create custom explicit web proxy services.

Web proxy services are similar to standard firewall services. You can configure web proxy services to define one or more protocols and port numbers that are associated with each web proxy service. Web proxy services can also be grouped into web proxy service groups.

One way in which web proxy services differ from firewall services is the protocol type you can select. The following protocol types are available:

- ALL
- CONNECT
- FTP
- HTTP

- SOCKS-TCP
- SOCKS-UDP

To add a web proxy service go to **Policy & Objects > Services** and select **Create New**. Set **Service Type** to **Explicit Proxy** and configure the service as required.

To add a web proxy service from the CLI enter:

```
config firewall service custom
   edit my-socks-service
      set explicit-proxy enable
      set category Web Proxy
      set protocol SOCKS-TCP
      set tcp-portrange 3450-3490
   end
```

To add a web proxy service group go to **Policy & Objects > Services** and select **Create New > Service Group**. Set **Type** to **Explicit Proxy** and add web proxy services to the group as required.

To add a web proxy service group from the CLI enter:

```
config firewall service group
   edit web-group
      set explicit-proxy enable
      set member webproxy my-socks-service
   end
```

Explicit web proxy firewall address URL patterns

You can add URL pattern addresses and address groups to control the destination URLs that explicit proxy users can connect to. To add a URL pattern to go to **Policy & Objects > Addresses**, select **Create New** and set the Type to **URL Pattern (Explicit Proxy)**. Add a URL or URL pattern that defines the URL or URLs that explicit proxy users should be limited to. Set the **Interface** to **any**.

For example to limit access to a single website:

www.fortinet.com

To limit access to websites from the same domain:

google.com

To limit access to a part of a website:

www.apple.com/ipad/

To add a URL pattern group, create several URL pattern addresses then go to **Policy & Objects > Addresses**, select **Create New > Group** and add URL patterns to the address group.

Then when creating explicit web proxy policies, select the URL pattern addresses or groups as the destination address.

URL patterns and HTTPS scanning

For HTTPS traffic, URL patterns can only be matched up to the root path. For example, consider the following URL pattern:

www.apple.com/ipad/

If a proxy user browses using HTTP, this URL pattern limits their access the iPad pages of www.apple.com. However, if a proxy user browses using HTTPS, they will be able to access all pages on www.apple.com.

Changing HTTP headers

You can create explicit web proxy profiles that can add, remove and change HTTP headers. The explicit web proxy profile can be added to a web explicit proxy policy and will be applied to all of the HTTP traffic accepted by that policy.

You can change the following HTTP headers:

- client-ip
- via header for forwarded requests
- via header for forwarded responses
- x-forwarded-for
- front-end-https

For each of these headers you can set the action to:

- Pass to forward the traffic without changing the header
- Add to add the header
- Remove to remove the header

You can also configure how the explicit web proxy handles custom headers. The proxy can add or remove custom headers from requests or responses. If you are adding a header you can specify the content to be included in the added header.

Create web proxy profiles from the CLI:

```
config web-proxy profile
   edit <name>
      set header-client-ip {add | pass | remove}
      set header-via-request {add | pass | remove}
      set header-via-response {add | pass | remove}
      set header-x-forwarded-for {add | pass | remove}
      set header-front-end-https {add | pass | remove}
         config headers
            edit <id>
               set action {add-to-request | add-to-response | remove-from-request |
                   remove-from-response}
               set content <string>
               set name <name>
            end
   end
```

Use the following command to add a web proxy profile to an explicit proxy policy:

```
config firewall explicit-proxy-policy
   edit <id>
      set webproxy-profile <name>
   end
```

Preventing the explicit web proxy from changing source addresses

By default in NAT/Route mode the explicit web proxy changes the source address of packets leaving the FortiGate to the IP address of the FortiGate interface that the packets are exiting from. In Transparent mode the source address is changed to the management IP.

This configuration hides the IP addresses of clients and allows packets to return to the FortiGate unit interface without having to route packets from clients. You can use the following command to configure the explicit web proxy to keep the original client's source IP address:

```
config firewall explicit-proxy-policy
   edit 0
      set proxy web
      set transparent enable
   end
```

Example users on an internal network browsing the Internet through the explicit web proxy with web caching, RADIUS authentication, web filtering, and virus scanning

This example describes how to configure the explicit web proxy for the example network shown below. In this example, users on the internal network connect to the explicit web proxy through the Internal interface of the FortiGate unit. The explicit web proxy is configured to use port 8888 so users must configure their web browser proxy settings to use port 8888 and IP address 10.31.101.100.

Example explicit web proxy network topology

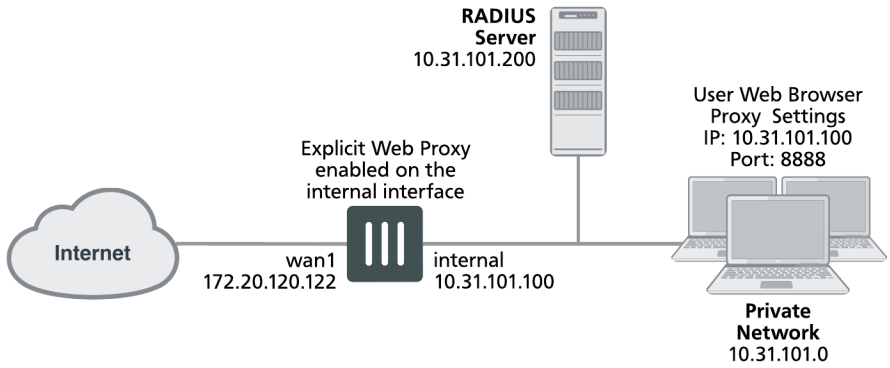

Explicit web proxy users must authenticate with a RADIUS server before getting access to the proxy. The explicit proxy policy that accepts explicit web proxy traffic applies per session authentication and includes a RADIUS server user group. The authentication rule also applies web filtering and virus scanning.

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Enable the explicit web proxy for HTTP and HTTPS and change the HTTP and HTTPS ports to 8888.
2. Enable the explicit web proxy on the internal interface.
3. Add a RADIUS server and user group for the explicit web proxy.

4. Add an authentication explicit proxy policy. Enable web caching. Add an authentication rule and enable antivirus and web filtering.

Configuring the explicit web proxy - web-based manager

Use the following steps to configure the explicit web proxy.

To enable and configure the explicit web proxy

1. Go to **System > Feature Select** and turn on the **Explicit Proxy** feature.
2. Go to **Network > Explicit Proxy** and change the following settings:

Enable Explicit Web Proxy	Select **HTTP/HTTPS**.
Listen on Interfaces	No change. This field will eventually show that the explicit web proxy is enabled for the Internal interface.
HTTP Port	8888
HTTPS Port	0
Realm	You are authenticating with the explicit web proxy.
Default Firewall Policy Action	Deny

3. Select **Apply**.

To enable the explicit web proxy on the Internal interface

1. Go to **Network > Interfaces**.
2. Edit the internal interface.
3. Select **Enable Explicit Web Proxy.**
4. Select **OK**.

To add a RADIUS server and user group for the explicit web proxy

1. Go to **User & Device > RADIUS Servers** and select **Create New** to add a new RADIUS server:

Name	RADIUS_1
Primary Server Name/IP	10.31.101.200
Primary Server Secret	RADIUS_server_secret

2. Select **OK**.
3. Go to **User & Device > User Groups** and select **Create New** to add a new user group.

Name	Explict_proxy_user_group

Type	Firewall
Remote Groups	RADIUS_1
Group Name	Any

4. Select **OK**.

To add an explicit proxy policy

1. Go to **Policy & Objects > Addresses** and select **Create New**.
2. Add a firewall address for the internal network:

Category	Address
Name	Internal_subnet
Type	Subnet / IP Range
Subnet / IP Range	10.31.101.0
Interface	Any

3. Go to **Policy & Objects > Explicit Proxy Policy** and select **Create New**.
4. Configure the explicit web proxy policy.

Explicit Proxy Type	Web
Source Address	Internal_subnet
Outgoing Interface	wan1
Destination Address	all
Action	AUTHENTICATE

5. Under **Configure Authentication Rules** select **Create New** to add an authentication rule:

Groups	Explicit_policy
Source User(s)	Leave blank
Schedule	always

6. Turn on **Antivirus** and **Web Filter** and select the **default** profiles for both.
7. Select the **default** proxy options profile.
8. Select **OK**.
9. Make sure **Enable IP Based Authentication** is not selected.
10. Turn on **Web Cache**.
11. Select **OK**.

Configuring the explicit web proxy - CLI

Use the following steps to configure the example explicit web proxy configuration from the CLI.

To enable the explicit web proxy on the Internal interface

1. Enter the following command to enable the explicit web proxy on the internal interface.

```
config system interface
   edit internal
      set explicit-web-proxy enable
end
```

To enable and configure the explicit web proxy

1. Enter the following command to enable the explicit web proxy and set the TCP port that proxy accepts HTTP and HTTPS connections on to 8888.

```
config web-proxy explicit
   set status enable
   set http-incoming-port 8888
   set https-incoming-port 8888
   set realm "You are authenticating with the explicit web proxy"
   set sec-default-action deny
end
```

To add a RADIUS server and user group for the explicit web proxy

1. Enter the following command to add a RADIUS server:

```
config user radius
   edit RADIUS_1
      set server 10.31.101.200
      set secret RADIUS_server_secret
end
```

2. Enter the following command to add a user group for the RADIUS server.

```
config user group
   edit Explicit_proxy_user_group
      set group-type firewall
      set member RADIUS_1
end
```

To add a security policy for the explicit web proxy

1. Enter the following command to add a firewall address for the internal subnet:

```
config firewall address
   edit Internal_subnet
      set type iprange
      set start-ip 10.31.101.1
      set end-ip 10.31.101.255
end
```

2. Enter the following command to add the explicit web proxy security policy:

```
config firewall explicit-proxy-policy
   edit 0
      set proxy web
```

```
        set dstintf wan1
        set srcaddr Internal_subnet
        set dstaddr all
        set action accept
        set service webproxy
        set webcache enable
        set identity-based enable
        set ipbased disable
        set active-auth-method basic
           config identity-based-policy
              edit 0
                 set groups Explicit_Proxy_user_group
                 set schedule always
                 set utm-status enable
                 set av-profile default
                 set webfilter-profile default
                 set profile-protocol-options default
              end
end
```

Testing and troubleshooting the configuration

You can use the following steps to verify that the explicit web proxy configuration is working as expected:

To test the explicit web proxy configuration

1. Configure a web browser on the internal subnet to use a web proxy server at IP address 10.31.101.100 and port 8888.
2. Browse to an Internet web page.
   The web browser should pop up an authentication window that includes the phrase that you added to the Realm option.
3. Enter the username and password for an account on the RADIUS server.
   If the account is valid you should be allowed to browse web pages on the Internet.
4. Close the browser and clear its cache and cookies.
5. Restart the browser and connect to the Internet.
   You could also start a second web browser on the same PC. Or you could start a new instance of the same browser as long as the browser asks for a user name and password again.

   You should have to authenticate again because identity-based policies are set to session-based authentication.
6. If this basic functionality does not work, check your FortiGate and web browser configuration settings.
7. Browse to a URL on the URL filter list and confirm that the web page is blocked.
8. Browse to http://eicar.org and attempt to download an anti-malware test file.
   The antivirus configuration should block the file.
   Sessions for web-proxy security policies do not appear on the Top Sessions dashboard widget and the count column for security policies does not display a count for explicit web proxy security policies.
9. You can use the following command to display explicit web proxy sessions

```
get test wad 60
IP based users:

Session based users:
   user:0x9c20778, username:User1, vf_id:0, ref_cnt:9
```

```
Total allocated user:1

Total user count:3, shared user quota:50, shared user count:3
```

This command output shows one explicit proxy user with user name `User1` authenticated using session-based authentication.

Explicit web proxy sessions and user limits

Web browsers and web servers open and close multiple sessions with the explicit web proxy. Some sessions open and close very quickly. HTTP 1.1 keepalive sessions are persistent and can remain open for long periods of time. Sessions can remain on the explicit web proxy session list after a user has stopped using the proxy (and has, for example, closed their browser). If an explicit web proxy session is idle for more than 3600 seconds it is torn down by the explicit web proxy. See RFC 2616 for information about HTTP keepalive/persistent HTTP sessions.

This section describes proxy sessions and user limits for both the explicit web proxy and the explicit FTP proxy. Session and user limits for the two proxies are counted and calculated together. However, in most cases if both proxies are active there will be many more web proxy sessions than FTP proxy sessions.

The FortiGate unit adds two sessions to its session table for every explicit proxy session started by a web browser and every FTP session started by an FTP client. An entry is added to the session table for the session from the web browser or client to the explicit proxy. All of these sessions have the same destination port as the explicit web proxy port (usually 8080 for HTTP and 21 for FTP). An entry is also added to the session table for the session between the exiting FortiGate interface and the web or FTP server destination of the session. All of these sessions have a FortiGate interface IP address and the source address of the session and usually have a destination port of 80 for HTTP and 21 for FTP.

Proxy sessions that appear in FortiView do not include the Policy ID of the web-proxy or ftp-proxy security policy that accepted them. However, the explicit proxy sessions include a destination port that matches the explicit proxy port number (usually 8080 for the web proxy and 21 for the FTP proxy). The proxied sessions from the FortiGate unit have their source address set to the IP address of the FortiGate unit interface that the sessions use to connect to their destinations (for example, for connections to the Internet the source address would be the IP address of the FortiGate interface connected to the Internet).

FortiOS limits the number of explicit proxy users. This includes both explicit FTP proxy and explicit web proxy users. The number of users varies by FortiGate model from as low as 10 to up to 18000 for high end models. You cannot raise this limit.

If your FortiGate unit is configured for multiple VDOMs you can go to **System > Global Resources** to view the maximum number of **Concurrent explicit proxy users** and optionally reduce the limit. You can also use the following command:

```
config global
   config system resource-limits
      set proxy 50
   end
end
```

To limit the number of explicit proxy users for a VDOM, from the web-based manager enable multiple VDOMs and go to **System > VDOM** and edit a VDOM or use the following command to change the number of explicit web proxy users for VDOM_1:

```
config global
   config system vdom-property
      edit VDOM_1
```

```
            set proxy 25
        end
    end
```

You can use the `diagnose wad user list` command to view the number of explicit web proxy users. Users may be displayed with this command even if they are no longer actively using the proxy. All idle sessions time out after 3600 seconds.

You can use the command `diagnose wad user clear` to clear current explicit proxy users. You can also use the command `diagnose wad user clear <user-name>` to clear individual users. This means delete information about all users and force them re-authenticate.

Users that authenticate with explicit web-proxy or ftp-proxy security policies do not appear in the **Monitor > Firewall Monitor** list and selecting **De-authenticate All Users** has no effect on explicit proxy users.

How the number of concurrent explicit proxy users is determined depends on their authentication method:

- For session-based authenticated users, each authenticated user is counted as a single user. Since multiple users can have the same user name, the proxy attempts to identify users according to their authentication membership (based upon whether they were authenticated using RADIUS, LADAP, FSAE, local database etc.). If a user of one session has the same name and membership as a user of another session, the explicit proxy assumes this is one user.
- For IP Based authentication, or no authentication, or if no web-proxy security policy has been added, the source IP address is used to determine a user. All sessions from a single source address are assumed to be from the same user.

The explicit proxy does not limit the number of active sessions for each user. As a result the actual explicit proxy session count is usually much higher than the number of explicit web proxy users. If an excessive number of explicit web proxy sessions is compromising system performance you can limit the amount of users if the FortiGate unit is operating with multiple VDOMs.

The FortiGate explicit FTP proxy

You can use the FortiGate explicit FTP proxy to enable explicit FTP proxying on one or more FortiGate interfaces. The explicit web and FTP proxies can be operating at the same time on the same or on different FortiGate interfaces.

Explicit FTP proxies are configured for each VDOM when multiple VDOMs are enabled.

In most cases you would configure the explicit FTP proxy for users on a network by enabling the explicit FTP proxy on the FortiGate interface connected to that network. Users on the network would connect to and authenticate with the explicit FTP proxy before connecting to an FTP server. In this case the IP address of the explicit FTP proxy is the IP address of the FortiGate interface on which the explicit FTP proxy is enabled.

Enabling the explicit FTP proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address.

If the FortiGate unit is operating in Transparent mode, users would configure their browsers to use a proxy server with the FortiGate unit management IP address.

The FTP proxy receives FTP sessions to be proxied at FortiGate interfaces with the explicit FTP proxy enabled. The FTP proxy uses FortiGate routing to route sessions through the FortiGate unit to a destination interface. Before a session leaves the exiting interface, the explicit FTP proxy changes the source addresses of the session packets to the IP address of the exiting interface. When the FortiGate unit is operating in Transparent mode the explicit web proxy changes the source addresses to the management IP address.

Example explicit FTP proxy topology

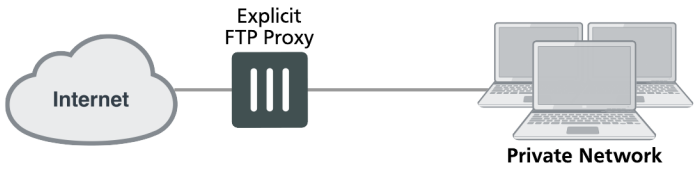

To allow anyone to anonymously log into explicit FTP proxy and connect to any FTP server you can set the explicit FTP proxy default firewall proxy action to accept. When you do this, users can log into the explicit FTP proxy with any username and password.

In most cases you would want to use explicit proxy policies to control explicit FTP proxy traffic and apply security features, access control/authentication, and logging. You can do this by keeping the default explicit FTP proxy firewall policy action to deny and then adding explicit FTP proxy policies. In most cases you would also want users to authenticate with the explicit FTP proxy. By default an anonymous FTP login is required. Usually you would add authentication to explicit FTP proxy policies. Users can then authenticate with the explicit FTP proxy according to users or user groups added to the policies. User groups added to explicit FTP proxy policies can use any authentication method supported by FortiOS including the local user database and RADIUS and other remote servers.

If you leave the default firewall policy action set to deny and add explicit FTP proxy policies, all connections to the explicit FTP proxy must match an or else they will be dropped. Sessions that are accepted are processed according to the ftp-proxy security policy settings.

You can also change the explicit FTP proxy default firewall policy action to accept and add explicit FTP proxy policies. If you do this, sessions that match explicit FTP proxy policies are processed according to the policy settings. Connections to the explicit FTP proxy that do not match an explicit FTP proxy policy are allowed and the users can authenticate with the proxy anonymously.

There are some limitations to the security features that can be applied to explicit FTP proxy sessions. See Security profiles, threat weight, device identification, and the explicit FTP proxy on page 126.

You cannot configure IPsec, SSL VPN, or Traffic shaping for explicit FTP proxy traffic. Explicit FTP proxy policies can only include firewall addresses not assigned to a FortiGate unit interface or with interface set to **any**. (On the web-based manager you must set the interface to **Any**. In the CLI you must `unset` the `associated-interface`.)

How to use the explicit FTP proxy to connect to an FTP server

To connect to an FTP server using the explicit FTP proxy, users must run an FTP client and connect to the IP address of a FortiGate interface on which the explicit FTP proxy is enabled. This connection attempt must use the configured explicit FTP proxy port number (default 21).

The explicit FTP proxy is not compatible with using a web browser as an FTP client. To use web browsers as FTP clients configure the explicit web proxy to accept FTP sessions.

The following steps occur when a user starts an FTP client to connect to an FTP server using the explicit FTP proxy. Any RFC-compliant FTP client can be used. This example describes using a command-line FTP client. Some FTP clients may require a custom FTP proxy connection script.

1. The user enters a command on the FTP client to connect to the explicit FTP proxy.

   For example, if the IP address of the FortiGate interface on which the explicit FTP proxy is enabled is 10.31.101.100, enter:

   ```
   ftp 10.31.101.100
   ```

2. The explicit FTP proxy responds with a welcome message and requests the user's FTP proxy user name and password and a username and address of the FTP server to connect to:

   ```
   Connected to 10.31.101.100.
   220 Welcome to Fortigate FTP proxy
   Name (10.31.101.100:user):
   ```

   You can change the message by editing the FTP Explicit Banner Message replacement message.

3. At the prompt the user enters their FTP proxy username and password and a username and address for the FTP server. The FTP server address can be a domain name or numeric IP address. This information is entered using the following syntax:

   ```
   <proxy-user>:<proxy-password>:<server-user>@<server-address>
   ```

   For example, if the proxy username and password are `p-name` and `p-pass` and a valid username for the FTP server is `s-name` and the server's IP address is `ftp.example.com` the syntax would be:

   ```
   p-name:p-pass:s-name@ftp.example.com
   ```

If the FTP proxy accepts anonymous logins `p-name` and `p-pass` can be any characters.

4. The FTP proxy forwards the connection request, including the user name, to the FTP server.
5. If the user name is valid for the FTP server it responds with a password request prompt.
6. The FTP proxy relays the password request to the FTP client.
7. The user enters the FTP server password and the client sends the password to the FTP proxy.
8. The FTP proxy relays the password to the FTP server.
9. The FTP server sends a login successful message to the FTP proxy.
10. The FTP proxy relays the login successful message to the FTP client.
11. The FTP client starts the FTP session.

All commands entered by the client are relayed by the proxy to the server. Replies from the server are relayed back to the FTP client.

Explicit FTP proxy session

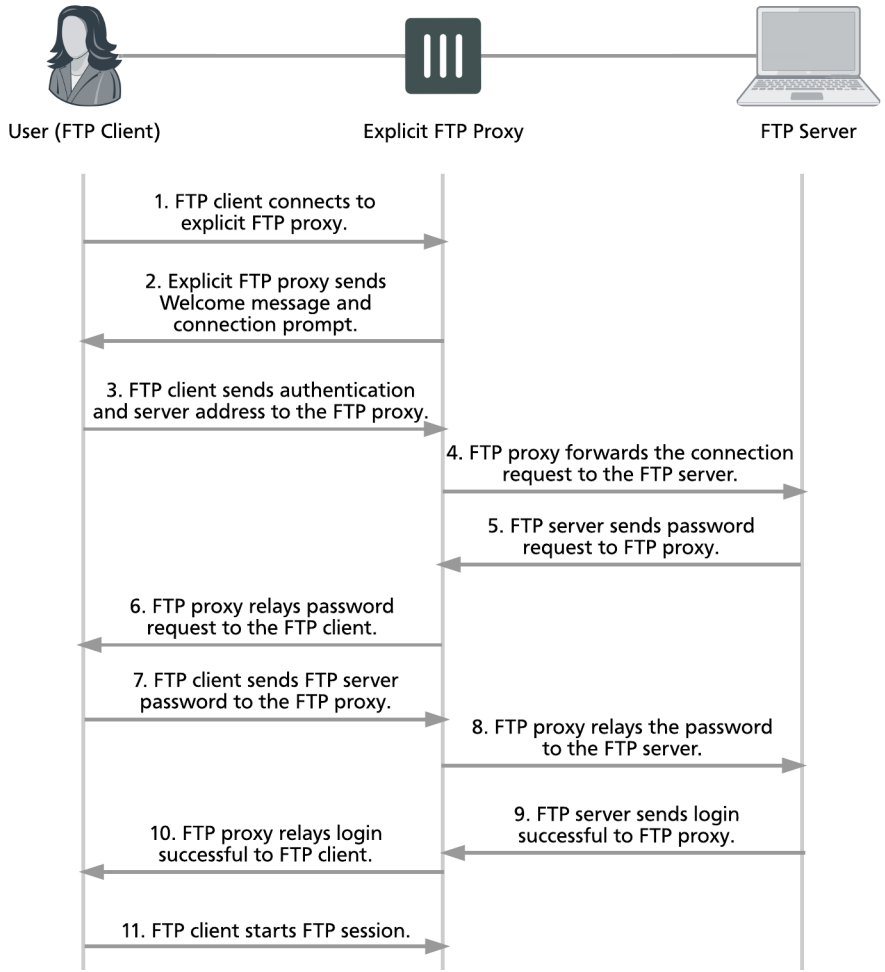

From a simple command line FTP client connecting to an the previous sequence could appear as follows:

```
ftp 10.31.101.100 21
Connected to 10.31.101.100.
220 Welcome to Fortigate FTP proxy
Name (10.31.101.100:user): p-name:p-pass:s-name@ftp.example.com
331 Please specify the password.
Password: s-pass
230 Login successful.
Remote system type is UNIX
Using binary mode to transfer files.
ftp>
```

General explicit FTP proxy configuration steps

You can use the following general steps to configure the explicit FTP proxy.

To enable the explicit FTP proxy - web-based manager:

1. Go to **Network > Explicit Proxy > Explicit FTP Proxy Options**. Select **Enable Explicit FTP Proxy** to turn on the explicit FTP proxy.
2. Select Apply.

   The **Default Firewall Policy Action** is set to **Deny** and requires you to add a explicit FTP proxy policy to allow access to the explicit FTP proxy. This configuration is recommended and is a best practice because you can use policies to control access to the explicit FTP proxy and also apply security features and authentication.
3. Go to **Network > Interfaces** and select one or more interfaces for which to enable the explicit web proxy. Edit the interface and select **Enable Explicit FTP Proxy**.

Enabling the explicit FTP proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address. If you enable the proxy on such an interface make sure authentication is required to use the proxy.

4. Go to **Policy & Objects > Explicit Proxy Policy** and select **Create New** and set the **Explicit Proxy Type** to **FTP.**

   You can add multiple explicit FTP proxy policies.
5. Configure the policy as required to accept the traffic that you want to be processed by the explicit FTP proxy.

   The source address of the policy should match client source IP addresses. The firewall address selected as the source address cannot be assigned to a FortiGate interface. The Interface field of the firewall address must be blank or it must be set to **Any**.

   The destination address of the policy should match the IP addresses of FTP servers that clients are connecting to. The destination address could be **all** to allow connections to any FTP server.

   If **Default Firewall Policy Action** is set to Deny, traffic sent to the explicit FTP proxy that is not accepted by an explicit FTP proxy policy is dropped. If **Default Firewall Policy Action** is set to Allow then all FTP proxy sessions that don't match a policy are allowed.

For example the following explicit FTP proxy policy allows users on an internal network to access FTP servers on the Internet through the wan1 interface of a FortiGate unit.

Explicit Proxy Type	FTP
Source Address	Internal_subnet
Outgoing Interface	wan1
Destination Address	all
Schedule	always
Action	ACCEPT

The following explicit FTP proxy policy requires users on an internal network to authenticate with the FortiGate unit before accessing FTP servers on the Internet through the wan1 interface.

Explicit Proxy Type	FTP
Source Address	Internal_subnet
Outgoing Interface	wan1
Destination Address	all
Action	AUTHENTICATE

6. Select **Create New** to add an **Authentication Rule** and configure the rule as follows:

Groups	Proxy-Group
Source Users	(optional)
Schedule	always

7. Add security profiles as required and select **OK**.
8. You can add multiple authentication rules to apply different authentication for different user groups and users and also apply different security profiles and logging settings for different users.
9. Select **OK**.

To enable the explicit FTP proxy - CLI:

1. Enter the following command to turn on the explicit FTP proxy. This command also changes the explicit FTP proxy port to 2121.

```
config ftp-proxy explicit
   set status enable
   set incoming-port 2121
end
```

The default explicit FTP proxy configuration has `sec-default-action` set to `deny` and requires you to add a security policy to allow access to the explicit FTP proxy.

2. Enter the following command to enable the explicit FTP proxy for the internal interface.

```
config system interface
   edit internal
      set explicit-ftp-proxy enable
   end
end
```

3. Use the following command to add a firewall address that matches the source address of users who connect to the explicit FTP proxy.

```
config firewall address
   edit Internal_subnet
      set type iprange
      set start-ip 10.31.101.1
      set end-ip 10.31.101.255
   end
```

The source address for a ftp-proxy security policy cannot be assigned to a FortiGate unit interface.

4. Use the following command to add an explicit FTP proxy policy that allows all users on the internal subnet to use the explicit FTP proxy for connections through the wan1 interface to the Internet.

```
config firewall explicit-proxy-policy
   edit 0
      set proxy ftp
      set dstintf wan1
      set scraddr Internal_subnet
      set dstaddr all
      set action accept
      set schedule always
end
```

5. Use the following command to add an explicit FTP proxy policy that allows authenticated users on the internal subnet to use the explicit FTP proxy for connections through the wan1 interface to the Internet.

```
config firewall explicit-proxy-policy
   edit 0
      set proxy ftp
      set dstintf wan1
      set scraddr Internal_subnet
      set dstaddr Fortinet-web-sites
      set action accept
      set schedule always
      set identity-based enable
         config identity-based-policy
            edit 1
               set groups Proxy-group
               set schedule always
      end
end
```

Restricting the IP address of the explicit FTP proxy

You can use the following command to restrict access to the explicit FTP proxy using only one IP address. The IP address that you specify must be the IP address of an interface that the explicit FTP proxy is enabled on. You might want to use this option if the explicit FTP proxy is enabled on an interface with multiple IP addresses.

For example, to require uses to connect to the IP address 10.31.101.100 to connect to the explicit FTP proxy:

```
config ftp-proxy explicit
   set incoming-ip 10.31.101.100
end
```

Restricting the outgoing source IP address of the explicit FTP proxy

You can use the following command to restrict the source address of outgoing FTP proxy packets to a single IP address. The IP address that you specify must be the IP address of an interface that the explicit FTP proxy is enabled on. You might want to use this option if the explicit FTP proxy is enabled on an interface with multiple IP addresses.

For example, to restrict the outgoing packet source address to 172.20.120.100:

```
config ftp-proxy explicit
   set outgoing-ip 172.20.120.100
end
```

Security profiles, threat weight, device identification, and the explicit FTP proxy

You can apply antivirus, data leak prevention (DLP), and SSL/SSH inspection to explicit FTP proxy sessions. Security profiles are applied by selecting them in an explicit FTP proxy policy or an authentication rule in an FTP proxy security policy.

Traffic accepted by explicit FTP proxy policies contributes to threat weight data.

The explicit FTP proxy is not compatible with device identification.

Explicit FTP proxy options and SSL/SSH inspection

Since the traffic accepted by the explicit FTP proxy is known to be FTP and since the ports are already known by the proxy, the explicit FTP proxy does not use the FTP port proxy options settings.

When adding UTM features to an FTP proxy security policy, you must select a proxy options profile. In most cases you can select the default proxy options profile. You could also create a custom proxy options profile.

The explicit FTP proxy supports the following proxy options:

- Block Oversized File and oversized file limit

The explicit FTP proxy does not support the following protocol options:

- Client comforting

Explicit FTP proxy sessions and antivirus

For explicit FTP proxy sessions, the FortiGate unit applies antivirus scanning to FTP file GET and PUT requests. The FortiGate unit starts virus scanning a file in an FTP session when it receives a file in the body of an FTP request.

Flow-based virus scanning is not available for explicit FTP proxy sessions. Even if the FortiGate unit is configured to use flow-based antivirus, explicit FTP proxy sessions use the regular virus database.

Example users on an internal network connecting to FTP servers on the Internet through the explicit FTP with RADIUS authentication and virus scanning

This example describes how to configure the explicit FTP proxy for the example network shown below. In this example, users on the internal network connect to the explicit FTP proxy through the Internal interface with IP address 10.31.101.100. The explicit web proxy is configured to use port 2121 so to connect to an FTP server on the Internet users must first connect to the explicit FTP proxy using IP address 10.31.101.100 and port 2121.

Example explicit FTP proxy network topology

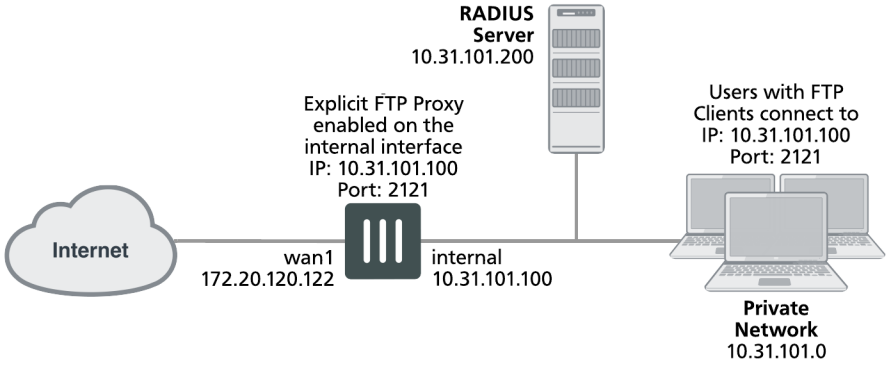

In this example, explicit FTP proxy users must authenticate with a RADIUS server before getting access to the proxy. To apply authentication, the security policy that accepts explicit FTP proxy traffic includes an identity based policy that applies per session authentication to explicit FTP proxy users and includes a user group with the RADIUS server in it. The identity based policy also applies UTM virus scanning and DLP.

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Enable the explicit FTP proxy and change the FTP port to 2121.
2. Enable the explicit FTP proxy on the internal interface.
3. Add a RADIUS server and user group for the explicit FTP proxy.
4. Add a user identity security policy for the explicit FTP proxy.
5. Enable antivirus and DLP features for the identity-based policy.

Configuring the explicit FTP proxy - web-based manager

Use the following steps to configure the explicit FTP proxy from FortiGate web-based manager.

To enable and configure the explicit FTP proxy

1. Go to **Network > Explicit Proxy > Explicit FTP Proxy Options** and change the following settings:

Enable Explicit FTP Proxy	Select.
Listen on Interface	No change. This field will eventually show that the explicit web proxy is enabled for the Internal interface.
FTP Port	2121
Default Firewall Policy Action	Deny

2. Select **Apply**.

To enable the explicit FTP proxy on the Internal interface

1. Go to **Network > Interfaces**, edit the Internal interface and select **Enable Explicit FTP Proxy**.

To add a RADIUS server and user group for the explicit FTP proxy

1. Go to **User & Device > RADIUS Servers**.
2. Select **Create New** to add a new RADIUS server:

Name	RADIUS_1
Primary Server Name/IP	10.31.101.200
Primary Server Secret	RADIUS_server_secret

3. Go to **User > User > User Groups** and select **Create New**.

Name	Explict_proxy_user_group
Type	Firewall
Remote groups	RADIUS_1
Group Name	ANY

4. Select **OK**.

To add a security policy for the explicit FTP proxy

1. Go to **Policy & Objects > Addresses** and select **Create New**.
2. Add a firewall address for the internal network:

Address Name	Internal_subnet
Type	Subnet
Subnet / IP Range	10.31.101.0
Interface	Any

3. Go to **Policy & Objects > Explicit Proxy Policy** and select **Create New**.
4. Configure the explicit FTP proxy security policy.

Explicit Proxy Type	FTP
Source Address	Internal_subnet
Outgoing Interface	wan1
Destination Address	all
Action	AUTHENTICATE

5. Under **Configure Authentication Rules** select **Create New** to add an authentication rule:

Groups	Explicit_policy
Users	Leave blank
Schedule	always

6. Turn on **Antivirus** and **Web Filter** and select the **default** profiles for both.
7. Select the **default** proxy options profile.
8. Select **OK**.
9. Make sure **Enable IP Based Authentication** is not selected and **Default Authentication Method** is set to **Basic**.
10. Select **OK**.

Configuring the explicit FTP proxy - CLI

Use the following steps to configure the example explicit web proxy configuration from the CLI.

To enable and configure the explicit FTP proxy

1. Enter the following command to enable the explicit FTP proxy and set the TCP port that proxy accepts FTP connections on to 2121.

```
config ftp-proxy explicit
   set status enable
   set incoming-port 2121
   set sec-default-action deny
end
```

To enable the explicit FTP proxy on the Internal interface

1. Enter the following command to enable the explicit FTP proxy on the internal interface.

```
config system interface
   edit internal
      set explicit-ftp-proxy enable
end
```

To add a RADIUS server and user group for the explicit FTP proxy

1. Enter the following command to add a RADIUS server:

```
config user radius
   edit RADIUS_1
      set server 10.31.101.200
      set secret RADIUS_server_secret
end
```

2. Enter the following command to add a user group for the RADIUS server.

```
config user group
   edit Explicit_proxy_user_group
      set group-type firewall
      set member RADIUS_1
end
```

To add a security policy for the explicit FTP proxy

1. Enter the following command to add a firewall address for the internal subnet:

```
config firewall address
   edit Internal_subnet
      set type iprange
      set start-ip 10.31.101.1
      set end-ip 10.31.101.255
end
```

2. Enter the following command to add the explicit FTP proxy security policy:

```
config firewall explicit-proxy-policy
   edit 0
      set proxy ftp
      set dstintf wan1
      set srcaddr Internal_subnet
      set dstaddr all
      set action accept
      set identity-based enable
      set ipbased disable
      set active-auth-method basic
         config identity-based-policy
            edit 0
               set groups Explicit_Proxy_user_group
               set schedule always
               set utm-status enable
               set av-profile default
               set profile-protocol-options default
            end
   end
```

Testing and troubleshooting the configuration

You can use the following steps to verify that the explicit FTP proxy configuration is working as expected. These steps use a command line FTP client.

To test the explicit web proxy configuration

1. From a system on the internal network start an FTP client and enter the following command to connect to the FTP proxy:

   ```
   ftp 10.31.101.100
   ```

   The explicit FTP proxy should respond with a message similar to the following:

   ```
   Connected to 10.31.101.100.
   220 Welcome to Fortigate FTP proxy
   Name (10.31.101.100:user):
   ```

2. At the prompt enter a valid username and password for the RADIUS server followed by a user name for an FTP server on the Internet and the address of the FTP server. For example, if a valid username and password on the RADIUS server is ex_name and ex_pass and you attempt to connect to an FTP server at ftp.example.com with user name s_name, enter the following at the prompt:

   ```
   Name (10.31.101.100:user):ex_name:ex_pass:s_name@ftp.example.com
   ```

3. You should be prompted for the password for the account on the FTP server.
4. Enter the password and you should be able to connect to the FTP server.
5. Attempt to explore the FTP server file system and download or upload files.
6. To test UTM functionality, attempt to upload or download an ECAR test file. Or upload or download a tex file containing text that would be matched by the DLP sensor.

   For eicar test files, go to http://eicar.org.

Explicit FTP proxy sessions and user limits

FTP clients do not open large numbers of sessions with the explicit FTP proxy. Most sessions stay open for a short while depending on how long a user is connected to an FTP server and how large the file uploads or downloads are. So unless you have large numbers of FTP users, the explicit FTP proxy should not be adding large numbers of sessions to the session table.

Explicit FTP proxy sessions and user limits are combined with explicit web proxy session and user limits. For information about explicit proxy session and user limits, see Explicit web proxy sessions and user limits on page 118.

FortiGate WCCP

The Web Cache Communication Protocol (WCCP) can be used to provide web caching with load balancing and fault tolerance. In a WCCP configuration, a WCCP server receives HTTP requests from user's web browsers and redirects the requests to one or more WCCP clients. The clients either return cached content or request new content from the destination web servers before caching it and returning it to the server which in turn returns the content to the original requestor. If a WCCP configuration includes multiple WCCP clients, the WCCP server load balances traffic among the clients and can detect when a client fails and failover sessions to still operating clients. WCCP is described by the Web Cache Communication Protocol Internet draft.

The sessions that are cached by WCCP depend on the configuration of the WCCP clients. If the client is a FortiGate unit, you can configure the port numbers and protocol number of the sessions to be cached. For example, to cache HTTPS traffic on port 443 the WCCP client port must be set to 443 and protocol must be set to 6. If the WCCP client should also cache HTTPS traffic on port 993 the client ports option should include both port 443 and 993.

On a FortiGate unit, WCCP sessions are accepted by a security policy before being cached. If the security policy that accepts sessions that do not match the port and protocol settings in the WCCP clients the traffic is dropped.

WCCP is configured per-VDOM. A single VDOM can operate as a WCCP server or client (not both at the same time). FortiGate units are compatible with third-party WCCP clients and servers. If a FortiGate unit is operating as an Internet firewall for a private network, you can configure it to cache and serve some or all of the web traffic on the private network using WCCP by adding one or more WCCP clients, configuring WCCP server settings on the FortiGate unit and adding WCCP security policies that accept HTTP session from the private network.

FortiGate units support WCCPv1 and WCCPv2. A FortiGate unit in NAT/Route or transparent mode can operate as a WCCP server. To operate as a WCCP client a FortiGate unit must be in NAT/Route mode. FortiGate units communicate between WCCP servers and clients over UDP port 2048. This communication can be encapsulated in a GRE tunnel or just use layer 2 forwarding.

A WCCP server can also be called a WCCP router. A WCCP client can also be called a WCCP cache engine.

WCCP service groups, service numbers, service IDs and well known services

A FortiGate unit configured as a WCCP server or client can include multiple server or client configurations. Each of these configurations is called a WCCP service group. A service group consists of one or more WCCP servers (or routers) and one or more WCCP clients working together to cache a specific type of traffic. The service group configuration includes information about the type of traffic to be cached, the addresses of the WCCP clients and servers and other information about the service.

A service group is identified with a numeric WCCP service ID (or service number) in the range 0 to 255. All of the servers and clients in the same WCCP service group must have service group configurations with the same WCCP service ID.

The value of the service ID provides some information about the type of traffic to be cached by the service group. Service IDs in the range 0 to 50 are reserved for well known services. A well known service is any service that is defined by the WCCP standard as being well known. Since the service is well known, just the service ID is required to identify the traffic to be cached.

Even though the well known service ID range is 0 to 50, at this time only one well known service has been defined. Its service ID 0, which is used for caching HTTP (web) traffic.

So to configure WCCP to cache HTTP sessions you can add a service group to the WCCP router and WCCP clients with a service ID of 0. No other information about the type of traffic to cache needs to be added to the service group.

Since service IDs 1 to 50 are reserved for well know services and since these services are not defined yet, you should not add service groups with IDs in the range 1 to 50.

FortiOS does allow you to add service groups with IDs between 1 and 50. Since these service groups have not been assigned well known services, however, they will not cache any sessions. Service groups with IDs 51 to 255 allow you to set the port numbers and protocol number of the traffic to be cached. So you can use service groups with IDs 51 to 255 to cache different kinds of traffic based on port numbers and protocol number of the traffic. Service groups 1 to 50; however, do not allow you to set port numbers or protocol numbers so cannot be used to cache any traffic.

To cache traffic other than HTTP traffic you must add service groups with IDs in the range 51 to 255. These service group configurations must include the port numbers and protocol number of the traffic to be cached. It is the port and protocol number configuration in the service group that determines what traffic will be cached by WCCP.

Example WCCP server and client configuration for caching HTTP sessions (service ID = 0)

Enter the following command to add a WCCP service group to a WCCP server that caches HTTP sessions. The IP address of the server is 10.31.101.100 and the WCCP clients are on the 10.31.101.0 subnet. The service

ID of this service group is 0.

```
config system wccp
   edit 0
      set router-id 10.31.101.100
      set server-list 10.31.101.0 255.255.255.0
   end
```

Enter the following commands to configure a FortiGate unit to operate as a WCCP client and add a service group that configures the client to cache HTTP sessions. The IP address of the server is 10.31.101.100 and IP address of this WCCP clients is 10.31.101.1 subnet. The service ID of this service group is 0.

```
config system settings
   set wccp-cache-engine enable
end

config system wccp
   edit 0
      set cache-id 10.31.101.1
      set router-list 10.31.101.100
   end
```

You cannot enter the `wccp-cache-engine enable` command if you have already added a WCCP service group. When you enter this command an interface named w.<vdom_name> is added to the FortiGate configuration (for example w.root). All traffic redirected from a WCCP router is considered to be received at this interface of the FortiGate unit operating as a WCCP client. A default route to this interface with lowest priority is added.

Example WCCP server and client configuration for caching HTTPS sessions

Enter the following command to add a service group to a WCCP server that caches HTTPS content on port 443 and protocol 6. The IP address of the server is 10.31.101.100 and the WCCP clients are on the 10.31.101.0 subnet. The service ID of this service group is 80.

```
config system settings
   set wccp-cache-engine enable
end

config system wccp
   edit 80
      set router-id 10.31.101.100
      set server-list 10.31.101.0 255.255.255.0
      set ports 443
      set protocol 6
   end
```

Enter the following commands to configure a FortiGate unit to operate as a WCCP client and add a service group that configures client to cache HTTPS sessions on port 443 and protocol 6. The IP address of the server is 10.31.101.100 and IP address of this WCCP clients is 10.31.101.1 subnet. The service ID of this service group must be 80 to match the service ID added to the server.

```
config system settings
   set wccp-cache-engine enable
end

config system wccp
   edit 80
      set cache-id 10.31.101.1
      set router-list 10.31.101.100
      set ports 443
      set protocol 6
   end
```

Example WCCP server and client configuration for caching HTTP and HTTPS sessions

You could do this by configuring two WCCP service groups as described in the previous examples. Or you could use the following commands to configure one service group for both types of traffic. The example also caches HTTP sessions on port 8080.

Enter the following command to add a service group to a WCCP server that caches HTTP sessions on ports 80 and 8080 and HTTPS sessions on port 443. Both of these protocols use protocol number 6. The IP address of the server is 10.31.101.100 and the WCCP clients are on the 10.31.101.0 subnet. The service ID of this service group is 90.

```
config system wccp
   edit 90
```

```
      set router-id 10.31.101.100
      set server-list 10.31.101.0 255.255.255.0
      set ports 443 80 8080
      set protocol 6
   end
```

Enter the following commands to configure a FortiGate unit to operate as a WCCP client and add a service group that configures client to cache HTTP sessions on port 80 and 8080 and HTTPS sessions on port 443. The IP address of the server is 10.31.101.100 and IP address of this WCCP clients is 10.31.101.1 subnet. The service ID of this service group must be 90 to match the service ID added to the server.

```
config system settings
   set wccp-cache-engine enable
end
config system wccp
   edit 90
      set cache-id 10.31.101.1
      set router-list 10.31.101.100
      set ports 443 80 8080
      set protocol 6
   end
```

Other WCCP service group options

In addition to using WCCP service groups to define the types of traffic to be cached by WCCP the following options are available for servers and clients.

Server configuration options

The server configuration must include the `router-id`, which is the WCCP server IP address. This is the IP address of the interface that the server uses to communicate with WCCP clients.

The `group-address` is used for multicast WCCP configurations to specify the multicast addresses of the clients.

The `server-list` defines the IP addresses of the WCCP clients that the server can connect to. Often the server list can be the address of the subnet that contains the WCCP clients.

The `authentication` option enables or disables authentication for the WCCP service group. Authentication must be enabled on all servers and clients in a service group and members of the group must have the same `password`.

The `forward-method` option specifies the protocol used for communication between the server and clients. The default forwarding method is GRE encapsulation. If required by your network you can also select to use unencapsulated layer-2 packets instead of GRE or select any to allow both. The `return-method` allows you to specify the communication method from the client to the server. Both GRE and layer-2 are supported.

The `assignment-method` determines how the server load balances sessions to the clients if there are multiple clients. Load balancing can be done using hashing or masking.

Client configuration options

The client configuration includes the `cache-id` which is the IP address of the FortiGate interface of the client that communicates with WCCP server. The `router-list` option is the list of IP addresses of the WCCP servers in the WCCP service group.

The `ports` option lists the port numbers of the sessions to be cached by the client and the `protocol` sets the protocol number of the sessions to be cached. For TCP sessions the protocol is 6.

The `service-type` option can be auto, dynamic or standard. Usually you would not change this setting.

The client configuration also includes options to influence load balancing including the `primary-hash`, `priority`, `assignment-weight` and `assignment-bucket-format`.

WCCP configuration overview

To configure WCCP you must create a service group that includes WCCP servers and clients. WCCP servers intercept sessions to be cached (for example, sessions from users browsing the web from a private network). To intercept sessions to be cached the WCCP server must include a security policy that accepts sessions to be cached and WCCP must be enabled in this security policy.

The server must have an interface configured for WCCP communication with WCCP clients. That interface sends and receives encapsulated GRE traffic to and from WCCP clients. The server must also include a WCCP service group that includes a service ID and the addresses of the WCCP clients as well as other WCCP configuration options.

To use a FortiGate unit as a WCCP client, the FortiGate unit must be set to be a WCCP client (or cache engine). You must also configure an interface on the client for WCCP communication. The client sends and receives encapsulated GRE traffic to and from the WCCP server using this interface.

The client must also include a WCCP service group with a service ID that matches a service ID on the server. The client service group also includes the IP address of the servers in the service group and specifies the port numbers and protocol number of the sessions that will be cached on the client.

When the client receives sessions from the server on its WCCP interface, it either returns cached content over the WCCP interface or connects to the destination web servers using the appropriate interface depending on the client routing configuration. Content received from web servers is then cached by the client and returned to the WCCP server over the WCCP link. The server then returns the received content to the initial requesting user web browser.

Finally you may also need to configure routing on the server and client FortiGate units and additional security policies may have to be added to the server to accept sessions not cached by WCCP.

Example caching HTTP sessions on port 80 using WCCP

In this example configuration (shown below), a FortiGate unit with host name WCCP_srv is operating as an Internet firewall for a private network is also configured as a WCCP server. The port1 interface of WCCP_srv is connected to the Internet and the port2 interface is connected to the internal network.

All HTTP traffic on port 80 that is received at the port2 interface of WCCP_srv is accepted by a port2 to port1 security policy with WCCP enabled. All other traffic received at the port2 interface is allowed to connect to the Internet by adding a general port2 to port1 security policy below the HTTP on port 80 security policy.

A WCCP service group is added to WCCP_srv with a service ID of 0 for caching HTTP traffic on port 80. The port5 interface of WCCP_srv is configured for WCCP communication.

A second FortiGate unit with host name WCCP_client is operating as a WCCP client. The port1 interface of WCCP_client is connected to port5 of WCCP_srv and is configured for WCCP communication.

WCCP_client is configured to cache HTTP traffic because it also has a WCCP service group with a service ID of 0.

WCCP_client connects to the Internet through WCCP_srv. To allow this, a port5 to port1 security policy is added to WCCP_srv.

FortiGate WCCP server and client configuration

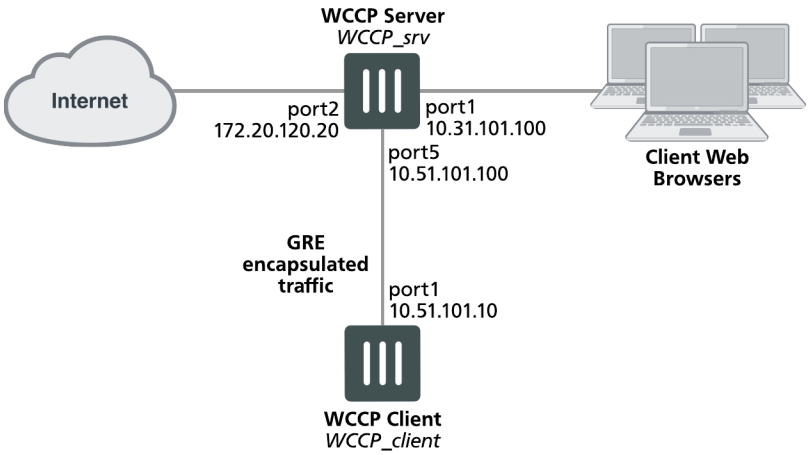

Configuring the WCCP server (WCCP_srv)

Use the following steps to configure WCCP_srv as the WCCP server for the example network. The example steps only describe the WCCP-related configuration.

To configure WCCP_srv as a WCCP server

1. Add a port2 to port1 security policy that accepts HTTP traffic on port 80 and is configured for WCCP:

```
config firewall policy
   edit 0
      set srtintf port2
      set dstintf port1
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service HTTP
      set wccp enable
      set nat enable
   end
```

2. Add another port2 to port1 security policy to allow all other traffic to connect to the Internet.

```
config firewall policy
   edit 0
      set srtintf port2
      set dstintf port1
      set srcaddr all
```

```
      set dstaddr all
      set action accept
      set schedule always
      set service ANY
      set nat enable
   end
```

3. Move this policy below the WCCP policy in the port2 to port1 policy list.
4. Enable WCCP on the port5 interface.

```
config system interface
   edit port5
      set wccp enable
   end
```

5. Add a WCCP service group with service ID 0.

```
config system wccp
   edit 0
      set router-id 10.51.101.100
      set server-list 10.51.101.0 255.255.255.0
   end
```

6. Add a firewall address and security policy to allow the WCCP_client to connect to the internet.

```
config firewall address
   edit WCCP_client_addr
      set subnet 10.51.101.10
end
config firewall policy
   edit 0
      set srtintf port5
      set dstintf port1
      set srcaddr WCCP_client_addr
      set dstaddr all
      set action accept
      set schedule always
      set service ANY
      set nat enable
   end
```

Configuring the WCCP client (WCCP_client)

Use the following steps to configure WCCP_client as the WCCP client for the example network. The example steps only describe the WCCP-related configuration.

To configure WCCP_client as a WCCP client

1. Configure WCCP_client to operate as a WCCP client.

```
config system settings
   set wccp-cache-engine enable
end
```

You cannot enter the `wccp-cache-engine enable` command if you have already added a WCCP service group. When you enter this command an interface named w.<vdom_name> is added to the FortiGate configuration (for example w.root). All traffic redirected from a WCCP router is considered to be received at this interface of the FortiGate unit operating as a WCCP client. A default route to this interface with lowest priority is added.

2. Enable WCCP on the port1 interface.

```
config system interface
   edit port1
      set wccp enable
end
```

3. Add a WCCP service group with service ID 0.

```
config system wccp
   edit 0
      set cache-id 10.51.101.10
      set router-list 10.51.101.100
end
```

Example caching HTTP sessions on port 80 and HTTPS sessions on port 443 using WCCP

This example configuration is the same as thatdescribed in Example caching HTTP sessions on port 80 using WCCP on page 136 except that WCCP now also cached HTTPS traffic on port 443. To cache HTTP and HTTPS traffic the WCCP service group must have a service ID in the range 51 to 255 and you must specify port 80 and 443 and protocol 6 in the service group configuration of the WCCP client.

Also the security policy on the WCCP_srv that accepts sessions from the internal network to be cached must accept HTTP and HTTPS sessions.

Configuring the WCCP server (WCCP_srv)

Use the following steps to configure WCCP_srv as the WCCP server for the example network. The example steps only describe the WCCP-related configuration.

To configure WCCP_srv as a WCCP server

1. Add a port2 to port1 security policy that accepts HTTP traffic on port 80 and HTTPS traffic on port 443 and is configured for WCCP:

```
config firewall policy
   edit 0
      set srtintf port2
      set dstintf port1
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service HTTP HTTPS
```

```
      set wccp enable
      set nat enable
   end
```

2. Add another port2 to port1 security policy to allow all other traffic to connect to the Internet.

```
config firewall policy
   edit 0
      set srtintf port2
      set dstintf port1
      set srcaddr all
      set dstaddr all
      set action accept
      set schedule always
      set service ANY

      set nat enable
   end
```

3. Move this policy below the WCCP policy in the port2 to port1 policy list.
4. Enable WCCP on the port5 interface.

```
config system interface
   edit port5
      set wccp enable
   end
```

5. Add a WCCP service group with service ID 90 (can be any number between 51 and 255).

```
config system wccp
   edit 90
      set router-id 10.51.101.100
      set server-list 10.51.101.0 255.255.255.0
   end
```

6. Add a firewall address and security policy to allow the WCCP_client to connect to the internet.

```
config firewall address
   edit WCCP_client_addr
      set subnet 10.51.101.10
   end
config firewall policy
   edit 0
      set srtintf port5
      set dstintf port1
      set srcaddr WCCP_client_addr
      set dstaddr all
      set action accept
      set schedule always
      set service ANY
      set nat enable
   end
```

Configuring the WCCP client (WCCP_client)

Use the following steps to configure WCCP_client as the WCCP client for the example network. The example steps only describe the WCCP-related configuration.

To configure WCCP_client as a WCCP client

1. Configure WCCP_client to operate as a WCCP client.

```
config system settings
   set wccp-cache-engine enable
end
```

You cannot enter the `wccp-cache-engine enable` command if you have already added a WCCP service group. When you enter this command an interface named w.<vdom_name> is added to the FortiGate configuration (for example w.root). All traffic redirected from a WCCP router is considered to be received at this interface of the FortiGate unit operating as a WCCP client. A default route to this interface with lowest priority is added.

2. Enable WCCP on the port1 interface.

```
config system interface
   edit port1
      set wccp enable
   end
```

3. Add a WCCP service group with service ID 90. This service group also specifies to cache sessions on ports 80 and 443 (for HTTP and HTTPS) and protocol number 6.

```
config system wccp
   edit 90
      set cache-id 10.51.101.10
      set router-list 10.51.101.100
      ports 80 443
      set protocol 6
   end
```

WCCP packet flow

The following packet flow sequence assumes you have configured a FortiGate unit to be a WCCP server and one or more FortiGate units to be WCCP clients.

1. A user's web browser sends a request for web content.
2. The FortiGate unit configured as a WCCP server includes a security policy that intercepts the request and forwards it to a WCCP client.

   The security policy can apply UTM features to traffic accepted by the policy.
3. The WCCP client receives the WCCP session.
4. The client either returns requested content to the WCCP server if it is already cached, or connects to the destination web server, receives and caches the content and then returns it to the WCCP server.
5. The WCCP server returns the requested content to the user's web browser.
6. The WCCP router returns the request to the client web browser.

   The client we browser is not aware that all this is taking place and does not have to be configured to use a web proxy.

Configuring the forward and return methods and adding authentication

The WCCP forwarding method determines how intercepted traffic is transmitted from the WCCP router to the WCCP cache engine. There are two different forwarding methods:

- GRE forwarding (the default) encapsulates the intercepted packet in an IP GRE header with a source IP address of the WCCP router and a destination IP address of the target WCCP cache engine. The results is a tunnel that allows the WCCP router to be multiple hops away from the WCCP cache server.
- L2 forwarding rewrites the destination MAC address of the intercepted packet to match the MAC address of the target WCCP cache engine. L2 forwarding requires that the WCCP router is Layer 2 adjacent to the WCCP client.

You can use the following command on a FortiGate unit configured as a WCCP router to change the forward and return methods to L2:

```
config system wccp
   edit 1
      set forward-method L2
      set return-method L2
end
```

You can also set the forward and return methods to any in order to match the cache server configuration.

By default the WCCP communication between the router and cache servers is unencrypted. If you are concerned about attackers sniffing the information in the WCCP stream you can use the following command to enable hash-based authentication of the WCCP traffic. You must enable authentication on the router and the cache engines and all must have the same password.

```
config system wccp
   edit 1
      set authentication enable
      set password <password>
end
```

WCCP Messages

When the WCCP service is active on a web cache server it periodically sends a WCCP HERE I AM broadcast or unicast message to the FortiGate unit operating as a WCCP router. This message contains the following information:

- Web cache identity (the IP address of the web cache server).
- Service info (the service group to join).

If the information received in the previous message matches what is expected, the FortiGate unit replies with a WCCP I SEE YOU message that contains the following details:

- Router identity (the FortiGate unit's IP address.
- Sent to IP (the web cache IP addresses to which the packets are addressed)

When both ends receive these two messages the connection is established, the service group is formed and the designated web cache is elected.

Troubleshooting WCCP

Two types of debug commands are available for debugging or troubleshooting a WCCP connection between a FortiGate unit operating as a WCCP router and its WCCP cache engines.

Real time debugging

The following commands can capture live WCCP messages:

```
diag debug en
diag debug application wccpd <debug level>
```

Application debugging

The following commands display information about WCCP operations:

```
get test wccpd <integer>
diag test application wccpd <integer>
```

Where `<integer>` is a value between 1 and 6:

1. Display WCCP stats
2. Display WCCP config
3. Display WCCP cache servers
4. Display WCCP services
5. Display WCCP assignment
6. Display WCCP cache status

Enter the following command to view debugging output:

```
diag test application wccpd 3
```

Sample output from a successful WCCP connection:

```
service-0 in vdom-root: num=1, usable=1
cache server ID:
len=44, addr=172.16.78.8, weight=4135, status=0
rcv_id=6547, usable=1, fm=1, nq=0, dev=3(k3),
to=192.168.11.55
ch_no=0, num_router=1:
192.168.11.55
```

Sample output from the same command from an unsuccessful WCCP connection (because of a service group password mismatch):

```
service-0 in vdom-root: num=0, usable=0
diag debug application wccpd -1
Sample output:
wccp_on_recv()-98: vdom-root recv: num=160, dev=3(3),
172.16.78.8->192.168.11.55
wccp2_receive_pkt()-1124: len=160, type=10, ver=0200,
length=152
wccp2_receive_pkt()-1150: found component:t=0, len=20
wccp2_receive_pkt()-1150: found component:t=1, len=24
wccp2_receive_pkt()-1150: found component:t=3, len=44
wccp2_receive_pkt()-1150: found component:t=5, len=20
wccp2_receive_pkt()-1150: found component:t=8, len=24
```

```
wccp2_check_security_info()-326: MD5 check failed
```

Diagnose commands

The following get and diagnose commands are available for troubleshooting WAN optimization, web cache, explicit proxy and WCCP.

get test {wad | wccpd} <test_level>

Display usage information about WAN optimization, explicit proxy, web cache, and WCCP applications. Use `<test_level>` to display different information.

```
get test wad <test_level>
get test wccpd <test_level>
```

Variable	Description
`wad`	Display information about WAN optimization, web caching, the explicit web proxy, and the explicit FTP proxy.
`wccpd`	Display information about the WCCP application.

Examples

Enter the following command to display WAN optimization tunnel protocol statistics. The http tunnel and tcp tunnel parts of the command output below shows that WAN optimization has been processing HTTP and TCP packets.

```
get test wad 1
WAD manager process status: pid=113 n_workers=1 ndebug_workers=0
```

Enter the following command to display all test options:

```
get test wad

WAD process 82 test usage:
   1: display process status
   2: display total memory usage.
   99: restart all WAD processes
   1000: List all WAD processes.
   1001: dispaly debug level name and values
   1002: dispaly status of WANOpt storages
   1068: Enable debug for all WAD workers.
   1069: Disable debug for all WAD workers.
   2yxx: Set No. xx process of type y as diagnosis process.
   3: display all fix-sized advanced memory stats
   4: display all fix-sized advanced memory stats in details
   500000..599999: cmem bucket stats (599999 for usage)
   800..899: mem_diag commands (800 for help & usage)
   800000..899999: mem_diag commands with 1 arg (800 for help & usage)
   80000000..89999999: mem_diag commands with 2 args (800 for help & usage)
   60: show debug stats.
```

```
61: discard all wad debug info that is currently pending
62xxx: set xxxM maximum ouput buffer size for WAD debug. 0, set back to default.
68: Enable process debug
69: Disable process debug
98: gracefully stopping WAD process
9xx: Set xx workers(0: default based on user configuration.)
```

diagnose wad

Display diagnostic information about the WAN optimization daemon (wad).

```
diagnose wad console-log {disable | enable)
diagnose wad debug-url {disable | enable)
diagnose wad filter {clear | dport | dst | list | negate | protocol | sport | src | vd}
diagnose wad history {clear | list}
diagnose wad session {clear | list}
diagnose wad stats {cache | cifs | clear | crypto | ftp | http | list | mapi | mem |
    scan | scripts | summary | tcp | tunnel}
diagnose wad user {clear | list}
diagnose wad tunnel {clear | list}1
diagnose wad webcache {clear | list} {10min | hour | day | 30days}
```

Variable	Description
`console-log`	Enable or disable displaying WAN optimization log messages on the CLI console.
`filter`	Set a filter for listing WAN optimization daemon sessions or tunnels. `clear` reset or clear the current log filter settings. `dport` enter the destination port range to filter by. `dst` enter the destination address range to filter by. `list` display the current log filter settings
`history`	Display statistics for one or more WAN optimization protocols for a specified period of time (the last 10 minutes, hour, day or 30 days).
`session`	Display diagnostics for WAN optimization sessions or clear active sessions.
`stats`	Display statistics for various parts of WAN optimization such as cache statistics, CIFS statistics, MAPI statistics, HTTP statistics, tunnel statistics etc. You can also clear WAN optimization statistics and display a summary.
`tunnel`	Display diagnostic information for one or all active WAN optimization tunnels. Clear all active tunnels. Clear all active tunnels.
`webcache`	Display web cache activity for the specified time period.

Example diagnose wad tunnel list

Enter the following command to list all of the running WAN optimization tunnels and display information about each one. The command output shows 10 tunnels all created by peer-to-peer WAN optimization rules (auto-detect set to off).

```
diagnose wad tunnel list

Tunnel: id=100 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web_servers id=100 ip=172.20.120.141
   SSL-secured-tunnel=no auth-grp=
   bytes_in=348 bytes_out=384

Tunnel: id=99 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web_servers id=99 ip=172.20.120.141
   SSL-secured-tunnel=no auth-grp=
   bytes_in=348 bytes_out=384

Tunnel: id=98 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web_servers id=98 ip=172.20.120.141
   SSL-secured-tunnel=no auth-grp=
   bytes_in=348 bytes_out=384

Tunnel: id=39 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web_servers id=39 ip=172.20.120.141
   SSL-secured-tunnel=no auth-grp=
   bytes_in=1068 bytes_out=1104

Tunnel: id=7 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web_servers id=7 ip=172.20.120.141
   SSL-secured-tunnel=no auth-grp=
   bytes_in=1228 bytes_out=1264

Tunnel: id=8 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web_servers id=8 ip=172.20.120.141
   SSL-secured-tunnel=no auth-grp=
   bytes_in=1228 bytes_out=1264

Tunnel: id=5 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web_servers id=5 ip=172.20.120.141
   SSL-secured-tunnel=no auth-grp=
   bytes_in=1228 bytes_out=1264

Tunnel: id=4 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web_servers id=4 ip=172.20.120.141
   SSL-secured-tunnel=no auth-grp=
   bytes_in=1228 bytes_out=1264
```

```
Tunnel: id=1 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web_servers id=1 ip=172.20.120.141
   SSL-secured-tunnel=no auth-grp=
   bytes_in=1228 bytes_out=1264

Tunnel: id=2 type=manual
   vd=0 shared=no uses=0 state=3
   peer name=Web_servers id=2 ip=172.20.120.141
   SSL-secured-tunnel=no auth-grp=
   bytes_in=1228 bytes_out=1264

Tunnels total=10 manual=10 auto=0
```

Example diagnose wad webcache list

This following command displays the web caching stats for the last 10 minutes of activity. The information displayed is divided into 20 slots and each slot contains stats for 30 seconds:

20 * 30 seconds = 600 seconds = 10 minutes

```
diagnose wad webcache list 10min
web cache history vd=0 period=last 10min
```

The first 20 slots are for HTTP requests in the last 10 minutes. Each slot of stats has four numbers, which is the total number of HTTP requests, the number of cacheable HTTP requests, the number of HTTP requests that are processed by the web cache (hits), and the number of HTTP requests that are processed without checking the web cache (bypass). There are many reasons that a HTTP request may bypass web cache.

```
total        cacheable     hits         bypass
------------ ------------- ------------ -------------
36           10            3            1
128          92            1            10
168          97            2            3
79           56            0            3
106          64            5            3
180          118           6            11
88           53            7            3
80           43            4            4
107          44            9            2
84           12            0            2
228          139           52           10
32           2             0            5
191          88            13           7
135          25            40           3
48           10            0            8
193          13            7            7
67           31            1            2
109          35            24           6
117          36            10           5
22           0             0            4
```

The following slots are for video requests in the last 10 minutes. Each slot has two numbers for each 30 seconds: total number of video requests, and the number of video requests that are processing using cached data.

```
video total  video hit
------------ -------------
0            0
0            0
0            0
0            0
0            0
0            0
0            0
0            0
0            0
0            0
0            0
0            0
0            0
0            0
0            0
```

The following 20 slots are for traffic details in last 10 minutes. Each slot has four numbers for 30 seconds each.

```
--- LAN ---                 --- WAN ---
bytes_in     bytes_out      bytes_in     bytes_out
------------ -------------- ------------ -------------
34360        150261         141086       32347
105408       861863         858501       100670
128359       1365919        1411849      127341
60103        602813         818075       59967
105867       1213192        1463736      97489
154961       1434784        1344911      158667
73967        370275         369847       70626
129327       602834         592399       123676
115719       663446         799445       111262
58151        724993         631721       59989
175681       2092925        1092556      166212
37805        33042          41528        37779
183686       1255118        1114646      172371
106125       904178         807152       81520
66147        473983         543507       66782
170451       1289530        1201639      165540
69196        544559         865370       68446
134142       579605         821430       132113
96895        668037         730633       89872
59576        248734         164002       59448
```

diagnose wacs

Display diagnostic information for the web cache database daemon (wacs).

```
diagnose wacs clear
diagnose wacs recents
diagnose wacs restart
diagnose wacs stats
```

Variable	Description
clear	Remove all entries from the web cache database.
recents	Display recent web cache database activity.
restart	Restart the web cache daemon and reset statistics.
stats	Display web cache statistics.

diagnose wadbd

Display diagnostic information for the WAN optimization database daemon (waddb).

```
diagnose wadbd {check | clear | recents | restart | stats}
```

Variable	Description
check	Check WAN optimization database integrity.
clear	Remove all entries from the WAN optimization database.
recents	Display recent WAN optimization database activity.
restart	Restart the WAN optimization daemon and reset statistics.
stats	Display WAN optimization statistics.

diagnose debug application {wad | wccpd} [<debug_level>]

View or set the debug level for displaying WAN optimization and web cache-related daemon debug messages. Include a <debug_level> to change the debug level. Leave the <debug_level> out to display the current debug level. Default debug level is 0.

```
diagnose debug application wad [<debug_level>]
diagnose debug application wccpd [<debug_level>]
```

Variable	Description
wad	Set the debug level for the WAN optimization daemon.
wccpd	Set the debug level for the WCCP daemon.

diagnose test application wad 2200

The debug level 2200 switches the debug to explicit proxy mode. You have to enter this debug level first. After that you have to type the command again with a different debug level to check the different explicit proxy statistics. To list what each debug level shows, follow these steps in any FortiGate device:

1. Enable explicit proxy globally and in one interface, to start the wad process. If the wad process is *not* running, you *cannot* list the options.
2. Once the wad process starts, type:

```
diagnose test application wad 2200
diagnose test application wad  ///// Do not type any debug level value to list all the options.
```

This is the output you will get:

```
# diagnose test application wad 2200
Set diagnosis process: type=wanopt index=0 pid=114
# diagnose test application wad
WAD process 114 test usage:
1: display process status
2: display total memory usage
99: restart all WAD processes
1000: List all WAD processes
1001: dispaly debug level name and values
1002: dispaly status of WANOpt storages
1068: Enable debug for all WAD workers
1069: Disable debug for all WAD workers
2yxx: Set No. xx process of type y as diagnosis process
3: display all fix-sized advanced memory stats
4: display all fix-sized advanced memory stats in details
500000..599999: cmem bucket stats (599999 for usage)
800..899: mem_diag commands (800 for help & usage)
800000..899999: mem_diag commands with 1 arg (800 for help & usage)
80000000..89999999: mem_diag commands with 2 args (800 for help & usage)
60: show debug stats
61: discard all wad debug info that is currently pending
62xxx: set xxxM maximum ouput buffer size for WAD debug (0: set back to default)
68: Enable process debug
69: Disable process debug
98: gracefully stopping WAD process
20: display all listeners
21: display TCP port info
22: display SSL stats
23: flush SSL stats
24: display SSL mem stats
70: display av memory usage
71xxxx: set xxxxMiB maximum AV memory (0: set back to default)
72: toggle av memory protection
73: toggle AV conserve mode (for debug purpose)
90: set to test disk failure
91: unset to test disk failure
92: trigger a disk failure event
100: display explicit proxy settings
101: display firewall policies
102: display security profile mapping for regular firewall policy
```

```
103: display Web proxy forwarding server and group
104: display DNS stats
105: display proxy redirection scan stats
106: list all used fqdns
107: list all firewall address
110: display current web proxy users
111: flush current web proxy users
112: display current web proxy user summary
113: display WAD fsso state
114: display HTTP digest stats
115: display URL patterns list of cache exemption or forward server
116: toggle dumping URL when daemon crashes
120: display Web Cache stats
121: flush Web Cache stats
122: flush idle Web cache objects
123: display web cache cache sessions
130: display ftpproxy stats
131: clear ftpproxy stats
132: list all current ftpproxy sessions
133: display all catched webfilter profiles
200: display WANopt profiles
201: display all peers
202: display video cache rules (patterns)
203: display all ssl servers
210: toggle disk-based byte-cache
211: toggle memory-based byte-cache
212: toggle cifs read-ahead
221: display tunnel protocol stats
222: flush tunnel protocol stats
223: display http protocol stats
224: flush http protocol stats
225: display cifs protocol stats
226: flush cifs protocol stats
227: display ftp protocol stats
228: flush ftp protocol stats
229: display mapi protocol stats
230: flush mapi protocol stats
231: display tcp protocol stats
232: flush tcp protocol stats
233: display all protocols stats
234: flush all protocols stats
240: display WAD tunnel stats
241: display tunnel compressor state
242: flush tunnel compressor stats
243: display Byte Cache DB state
244: flush Byte Cache DB stats
245: display Web Cache DB state
246: flush Web Cache DB stats
247: display cache state
248: flush cache stats
249: display memory cache state
250: flush memory cache stats
261yxxx: set xxx concurrent Web Cache session for object storage y
262yxxx: set xxxK(32K, 64K,...) unconfirmed write/read size per Web Cache object for
    object storage y
263yxxxx: set xxxxK maximum ouput buffer size for object storage y
```

```
264yxx: set lookup lowmark (only if more to define busy status) to be xx for object
    storage y
265yxxx: set xxxK maximum ouput buffer size for byte storage y
266yxxx: set number of buffered add requests to be xxx for byte storage y
267yxxxx: set number of buffered query requests to be xxxx for byte storage y
268yxxxxx: set number of concurrent query requests to be xxxxx for byte storage y
```

FORTINET®

High Performance Network Security